Unified computational framework for the efficient solution of n-field coupled problems with monolithic schemes*

Francesc Verdugo and Wolfgang A. Wall

Institute for Computational Mechanics
Technische Universität München
Boltzmannstr. 15
85748 Garching b. München
Email: {verdugo,wall}@lnm.mw.tum.de

arXiv:1605.01522v2 [math.NA] 14 Jul 2016

Abstract

In this paper, we propose and evaluate the performance of a unified computational framework for preconditioning systems of linear equations resulting from the solution of coupled problems with monolithic schemes. The framework is composed by promising application-specific preconditioners presented previously in the literature with the common feature that they are able to be implemented for a generic coupled problem, involving an arbitrary number of fields, and to be used to solve a variety of applications. The first selected preconditioner is based on a generic block Gauss-Seidel iteration for uncoupling the fields, and standard algebraic multigrid (AMG) methods for solving the resulting uncoupled problems. The second preconditioner is based on the semi-implicit method for pressure-linked equations (SIMPLE) which is extended here to deal with an arbitrary number of fields, and also results in uncoupled problems that can be solved with standard AMG. Finally, a more sophisticated preconditioner is considered which enforces the coupling at all AMG levels, in contrast to the other two techniques which resolve the coupling only at the finest level. Our purpose is to show that these methods perform satisfactory in quite different scenarios apart from their original applications. To this end, we consider three very different coupled problems: thermo-structure interaction, fluid-structure interaction and a complex model of the human lung. Numerical results show that these general purpose methods are efficient and scalable in this range of applications.

Keywords: Coupled problems · Monolithic solution schemes · Block preconditioners · Algebraic Multigrid · High Performance Computing

1 Introduction

The numerical simulation of coupled problems requires special solution strategies for solving the coupling between the underlying physical fields. The so-called *partitioned* methods [1] are often the preferred choice in industrial applications because they allow to reuse existing (black-box) solvers for the individual fields. However, this approach is unstable for many challenging strongly coupled problems [2, 3] and, therefore, another family of methods called *monolithic* are required in a variety of complex settings. It has also been shown that monolithic schemes are often preferable in terms of efficiency as compared to partitioned ones. For that reason, monolithic methods have been the preferred option for solving many strongly coupled problems, see e.g. [4, 5, 6, 7, 8, 9, 10, 11, 12, 13].

The price to be paid for the extra robustness of monolithic methods is a more challenging system of linear equations. The system matrix is a big sparse matrix with a special block structure representing each of the underlying physical fields and frequently has a very bad condition number. In real-world applications, iterative methods as GMRES [14] are required to attack this linear system, which requires efficient preconditioners for addressing the bad conditioning of the problem. Selecting a suitable preconditioner

*Article accepted for publication in *Computer Methods in Applied Mechanics and Engineering*, July 14 2016.

is the key point in the solution process and it is crucial to benefit from the robustness and efficiency of a monolithic approach.

The preconditioners proposed in the literature for coupled problems are usually complex methods designed for specific applications. This approach has the advantage that the preconditioner can be optimized for the particular problem leading to efficient solvers, but on the other hand, specific implementations have to be developed and maintained by experienced users for each particular application. This is specially an issue in "general purpose" multiphysics codes designed to simulate different kinds of problems, because specific implementations have to be developed and maintained separately.

In this paper, we explore an alternative approach. We select from the preconditioners available in the literature several techniques that are able to be implemented for a generic coupled problem, involving an arbitrary number of fields, and to be used to solve a variety of applications. We have implemented these methods in a unified computational framework allowing researchers to test and use them conveniently for different problem types without the need of developing application-specific implementations. Our purpose is to show that such a framework is realizable and that the selected methods are efficient and able to cope with a wide range of applications.

The first type of preconditioners proposed here is based on the so-called *block preconditioners*, see e.g. [11], which lead to efficient solvers by taking advantage of the block structure of the monolithic system. The key idea is to use approximate block inverses in order to untangle the coupling between the physical fields, and then, to use efficient algebraic multigrid (AMG) [15, 16, 17] solvers for the resulting uncoupled problems. This approach has been widely used in the literature for particular applications as fluid-structure interaction (FSI) [8, 12], magneto hydro dynamics [5, 6, 18] and thermo-structure interaction (TSI) [7], among others. In this context, we propose two different methods: a generalization of our previous works [7, 8] considering a block Gauss-Seidel (BGS) iteration for uncoupling the fields, and a generalization of [5] which considers an extension of the SIMPLE method [19] for the same purpose.

The drawback of standard block preconditioning based on BGS or SIMPLE is that the coupling is resolved only at the finest multigrid level. Thus, using efficient AMG solvers for the underlying problems does not necessarily imply a good treatment of the coupling and a fast global solution. This drawback is overcome by Gee et al. [8] who propose an enhanced block preconditioner (referred to as *monolithic AMG*) for FSI applications, which enforces the coupling at all multigrid levels. This often results in a better performance of the solver in the FSI problems studied in [8]. Monolithic AMG preconditioners are a promising approach for other coupled problems as well but, to our knowledge, this strategy has only been applied to FSI so far. One of the reasons why this approach has not been used and transferred to other problems so far is exactly due to the reasons given above: the required special know-how and cumbersome, necessary separate implementation. Here, we also propose an extension of the monolithic AMG preconditioner to a generic coupled problem. This is the first time that this preconditioner type is transferred to other problem types than the original FSI-application. Other more sophisticated block preconditioners (not discussed here) can also circumvent the degradation in performance associated with standard block preconditioners based on BGS or SIMPLE. For instance, the so-called pressure-convection-diffusion (PCD) or the least square commutator (LSC) preconditioners are successful alternatives in the context of the Navier-Stokes equations [20]. For magneto hydro dynamics (MHD), suitable alternatives are the preconditioners based on operator splitting or approximate commutators as shown in [6, 21] respectively.

The block preconditioners discussed in this paper have been implemented in the high performance multiphysics code BACI [22] and have been tested with three very different applications: TSI, FSI and a complex respiratory model of the human lung proposed by Yoshihara et al. [23, 24]. The numerical results show that these type of preconditioners show good performance for these different applications and suggest that the methods can be potentially applied to other coupled problems as well.

The block preconditioners studied here have been also implemented up to a certain extent in other related software packages and successfully applied to model problems different from those considered in this paper. One of the first object-oriented libraries supporting block preconditioners was already

presented back in 1998 [25]. This package included matrices with block structure and classical methods for uncoupling the fields as block Jacobi or block Gauss-Seidel. However, it did not support more advanced features as recursively nested block matrices which where not introduced until recently in modern software tools as Playa [26], Teko [27, 28] or PCFIELDSPLIT [29, 30]. Playa is a general high-level interface for manipulating block matrices which in particular allows the creation of block preconditioners. On the other hand, Teko is a block preconditioning environment that includes state-of-the-art preconditioners for Navier-Stokes [20] and MHD [21] ready to use, modify and extend by the users. Both Teko and Playa are part of the Trilinos project [31]. Finally, PCFIELDSPLIT is a component of the PETSc library [32] and shares some similarities with Teko like the creation and run-time design of complex block preconditioners including modern methods based on Schur complements.

The computational framework discussed in this paper shares many of the capabilities of Teko and PCFIELDSPLIT including the support for nested block matrices and run-time design of preconditioners. In contrast to these two packages, we have included and tested in our framework the advanced monolithic AMG solver where a multilevel hierarchy of block matrices has to be created. Moreover, we have provided the ability of combining monolithic AMG solvers with incomplete block factorizations based on the Schur complement as it is shown in Section 6 for a pulmonary mechanics example.

The remainder of the paper is structured as follows. In section 2, the monolithic approach is presented for the solution of a generic coupled problem leading to the system of linear equations to be solved with the proposed techniques. Section 3 includes a brief overview of solution techniques for linear systems and describes the role of the preconditioner. In section 4, the proposed general framework for preconditioning coupled problems is presented. In Section 5, some remarks on the implementation are given. Section 6 details the three applications considered for demonstrating the versatility of the methods and Section 7 contains the corresponding numerical examples. Finally, the paper is concluded in Section 8.

2 Monolithic system of linear equations

We focus on the solution of coupled problems that are described by partial differential equations that after the usual space and time discretization result in a set of non-linear algebraic equations to be solved at each time step. For this purpose, we consider a generic set of non-linear equations

$$\begin{aligned} \mathbf{f}_1(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N) &= \mathbf{0}, \\ \mathbf{f}_2(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N) &= \mathbf{0}, \\ &\vdots \\ \mathbf{f}_N(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N) &= \mathbf{0}, \end{aligned} \tag{1}$$

representing the discrete equations associated with a generic coupled problem with N coupled fields. The non-linear vector-valued functions $\mathbf{f}_1, \dots, \mathbf{f}_N$ represent the discrete versions of the state equations for each of the fields, and the vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N$ stand for the corresponding discrete solution vectors. Particular coupled problems leading to systems like (1) are given in Sections 6.1, 6.2 and 6.3 for the applications considered in the paper.

The discrete equations (1) form a fully-coupled non-linear problem which is solved here using a monolithic Newton method. To facilitate the discussion, equations (1) are written more compactly as

$$\mathbf{F}(\mathbf{U}) = \mathbf{0} \quad \text{with} \quad \mathbf{F} := \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \\ \vdots \\ \mathbf{f}_N \end{bmatrix} \quad \text{and} \quad \mathbf{U} := \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_N \end{bmatrix}. \tag{2}$$

With this notation, the monolithic Newton iteration for solving problem (1) is:

$$\mathbf{U}^{(i+1)} = \mathbf{U}^{(i)} + \Delta\mathbf{U}^{(i)} \tag{3}$$

where $\mathbf{U}^{(i)}$ stands for the full solution vector at the i-th iteration, and the correction $\Delta\mathbf{U}^{(i)}$ is obtained by solving the system of linear equations

$$\left[\frac{\partial \mathbf{F}}{\partial \mathbf{U}}\right]^{(i)} \Delta\mathbf{U}^{(i)} = -\mathbf{F}^{(i)}. \tag{4}$$

System (4) results from the usual linearization of the non-linear problem (1) at the i-th Newton step and it involves the linearization of all the field equations $\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_N$ with respect to all the variables $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_N$. This leads to a system matrix with an $N \times N$ block structure:

$$\begin{bmatrix} \frac{\partial \mathbf{f}_1}{\partial \mathbf{u}_1} & \frac{\partial \mathbf{f}_1}{\partial \mathbf{u}_2} & \cdots & \frac{\partial \mathbf{f}_1}{\partial \mathbf{u}_N} \\ \frac{\partial \mathbf{f}_2}{\partial \mathbf{u}_1} & \frac{\partial \mathbf{f}_2}{\partial \mathbf{u}_2} & \cdots & \frac{\partial \mathbf{f}_2}{\partial \mathbf{u}_N} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{f}_N}{\partial \mathbf{u}_1} & \frac{\partial \mathbf{f}_N}{\partial \mathbf{u}_2} & \cdots & \frac{\partial \mathbf{f}_N}{\partial \mathbf{u}_N} \end{bmatrix}^{(i)} \begin{bmatrix} \Delta\mathbf{u}_1 \\ \Delta\mathbf{u}_2 \\ \vdots \\ \Delta\mathbf{u}_N \end{bmatrix}^{(i)} = - \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \\ \vdots \\ \mathbf{f}_N \end{bmatrix}^{(i)} . \tag{5}$$

It is implicitly assumed that some of the blocks in (5) can also be zero.

The system (5) is one of the main difficulties associated with the solution of coupled problems with monolithic schemes. From the linear solver perspective, system (5) is particularly challenging because it involves all the physical fields. As a result, the system size is especially large and the condition number is usually very bad due to the different scaling magnitudes associated with the distinct nature of the underlying physics.

The goal of this paper is to propose solution strategies, i.e. preconditioners, for the efficient solution of system (5) for a wide range of coupled problems.

3 Iterative solution techniques

Before presenting the preconditioners for coupled problems, since this paper is targeting people dealing with modeling and simulation of coupled problems which are not necessarily solver experts, we briefly describe the linear solver used for solving the monolithic system (5), and which role the preconditioners play within the method. Moreover, we briefly comment on multigrid preconditioning techniques for single-field problems which are the main ingredient for building the proposed preconditioners for coupled problems. Obviously, it is beyond the scope of this paper to provide an adequate in-depth discussion of these topics. The idea is more to provide some intuition on the subject and some references for further consideration by the reader.

3.1 Krylov subspace methods

The resolution of the monolithic system (5) with direct methods is not affordable in many practical applications. The fill-in of the underlying sparse matrices leads to extreme memory requirements and computation times making direct methods inefficient. Moreover, direct methods are complex to implement in parallel and have a rather poor scalability. For this reason, the so-called Krylov subspace methods [33] are often preferred for solving large sparse systems like (5) in parallel computers.

Among the most known variants of Krylov subspace methods are the conjugate gradient method [34] for symmetric positive definite matrices, and the GMRES method [14] which is able to deal with generic

non-singular matrices. In the following, we consider the GMRES method implemented in the AztecOO package [35] for solving the coupled system (5). We choose GMRES because is the most general Krylov subspace method able to deal with non-symmetric matrices, but the same discussion applies for other Krylov solvers.

For a given system of linear equations

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \tag{6}$$

the GMRES method furnishes a sequence of approximate solutions $\mathbf{x}^k$ that converge to the exact solution $\mathbf{x}$ starting from an initial guess $\mathbf{x}^0$. The convergence rate strongly depends on the spectral properties of the system matrix $\mathbf{A}$. In the easiest case, when matrix $\mathbf{A}$ is symmetric definite, the convergence rate strongly depends on its condition number (see e.g. [36]):

$$\kappa_2(\mathbf{A}) := \frac{\sigma_{\max}(\mathbf{A})}{\sigma_{\min}(\mathbf{A})},$$

where $\sigma_{\max}(\mathbf{A})$ and $\sigma_{\min}(\mathbf{A})$ are the maximum and minimum eigenvalue of matrix $\mathbf{A}$ respectively. If the condition number is small, i.e. $\kappa_2(\mathbf{A}) \approx 1$, then, the convergence is expected to be fast. Whereas, for large condition numbers, the convergence is frequently very slow. The connection between the eigenvalues of $\mathbf{A}$ and the convergence behavior of the Krylov method becomes more involved when the matrix is not symmetric definite. The condition number is only a simplified convergence estimate in these situations. For a normal matrix $\mathbf{A}$, the convergence depends on how the eigenvalues are distributed on the complex plain. If the eigenvalues are clustered within a small region away from the origin, fast convergence of the method is expected, see e.g. [37]. Unfortunately, the convergence behavior of Krylov methods is still not fully understood for general nonsingular matrices which are not symmetric definite nor normal, see e.g. [38]. Consequently, the condition number can only be used as a simplified convergence estimate for non symmetric definite matrices.

The condition number of the monolithic matrix in (5) is usually very large due to several reasons including different scaling magnitudes associated with different underlying physics, strong hyperbolic and elliptic operators, and fine finite element meshes. Efficient preconditioners have to be used for improving the spectral properties of the system matrix in order use Krylov methods efficiently. Without the use of a preconditioner, Krylov methods do not converge at all for many practical applications. However, if efficient preconditioners are used, Krylov methods are among the most efficient techniques for solving this type of systems, see for instance [13].

3.2 The idea of preconditioning

A preconditioner $\mathbf{P}$ for a generic system (6) is a linear operator that tries to represent the action of the inverse of the system matrix $\mathbf{A}$, namely $\mathbf{P} \approx \mathbf{A}^{-1}$. In practice, the preconditioner is a procedure which takes a given vector $\bar{\mathbf{b}} \in \mathbb{R}^r$ (r is the number of rows or columns in $\mathbf{A}$) and returns a vector $\bar{\mathbf{x}} \in \mathbb{R}^r$ being the approximate solution of a system of equations formed by the system matrix $\mathbf{A}$ and taking vector $\bar{\mathbf{b}}$ as the right hand side:

$$\begin{array}{rccl} \mathbf{P}: & \mathbb{R}^r & \longrightarrow & \mathbb{R}^r \\ & \bar{\mathbf{b}} & \longmapsto & \bar{\mathbf{x}} \approx \mathbf{A}^{-1}\bar{\mathbf{b}} \end{array}. \tag{7}$$

The preconditioner can be seen as a cheap linear solver used to improve the convergence of a GMRES iteration. In fact, the preconditioner is frequently referred to as *the solver*, whereas the outer GMRES iteration as *the accelerator*. This is because the preconditioner usually requires more attention and problem-specific fine tuning than the GMRES iteration which is a generic procedure that can be considered as a black box.

For a given preconditioner $\mathbf{P}$, the preconditioned version of the system (6) is

$$(\mathbf{A}\mathbf{P})\left(\mathbf{P}^{-1}\mathbf{x}\right) = \mathbf{b}.$$

Solving the preconditioned system requires first solving

$$(\mathbf{AP})\,\mathbf{y} = \mathbf{b} \tag{8}$$

for the auxiliary variable $\mathbf{y} := \mathbf{P}^{-1}\mathbf{x}$, and finally obtaining the unknown $\mathbf{x}$ with an application of the preconditioning operator, $\mathbf{x} = \mathbf{P}\mathbf{y}$.

If the preconditioner $\mathbf{P}$ is a good approximation of $\mathbf{A}^{-1}$, then the matrix $\mathbf{AP}$ is expected to be close to the identity and the condition number of the preconitioned system (8) is expected to be $\kappa_2(\mathbf{AP}^{-1}) \approx 1$ which is usually associated with a good convergence of the GMRES method. The extra cost to be paid for this improved conditioning is an application of the preconditioner at each GMRES iteration, see [14].

Our purpose is to present different preconditioners (i.e. efficient solvers) for the monolithic system (5) to be used within a GMRES iteration for solving the system efficiently. The key idea of the proposed methods is to reuse existing preconditioning techniques that are known to be efficient for the underlying single-field problems. In this context, one of the most successful preconditioners are multigrid methods.

3.3 Multigrid methods

Multigrid methods [15] are efficient solution techniques designed for systems of linear equations resulting from the discretization of partial differential equations. The basic idea is to solve the system with an inexpensive iterative method (e.g. few Gauss-Seidel iterations), and then to approximate the unresolved error using a coarser version of the problem (that can be constructed by discretizing the underlying PDE using a coarser mesh). This idea can be applied recursively to an arbitrary number of coarser levels until the system matrix is small enough to compute an accurate coarse correction using a direct solver.

Defining a multigrid method requires introducing a hierarchy of coarser versions of the original discretization. In this context, one can distinguish two main approaches: either geometric multigrid methods [39], which explicitly consider coarser computational meshes, or algebraic multigrid methods (AMG) [15], which achieve coarser approximation spaces precluding the computation of new meshes by using information of the system matrix. The chosen coarsening procedure has little algorithmic relevance here. In the examples below, we use the so-called aggregation-based AMG methods [40] provided in the packages MueLu [41] and ML [42] because they are able to deal with complex 3D geometries.

The computed hierarchy of discretizations form the so-called multigrid levels. We denote the finest level as $\ell = 1$, whereas $\ell = \ell_{\max}$ corresponds to the coarsest level, see Figure 1. The main ingredients of a multigrid method are the following objects associated with each one of the levels $\ell = 1, \dots, \ell_{\max}$: 1) the projection matrices $\mathbf{P}^\ell$, that translate vectors form the coarse discretization on level $\ell + 1$ to the fine discretization on level ℓ, 2) Restriction matrices $\mathbf{R}^\ell$, that transforms vectors from level ℓ to level $\ell + 1$, and 3) iterative methods $\boldsymbol{S}^\ell$ (referred to as *the smoothers*) used to solve the system of equations associated with the discretization on level ℓ. Using the transfer operators, the matrices $\mathbf{A}^\ell$ associated with the discretization on levels $\ell = 2, \dots, \ell_{\max}$ are computed from the fine level matrix $\mathbf{A}^1$ as

$$\mathbf{A}^{\ell+1} = \mathbf{R}^\ell \mathbf{A}^\ell \mathbf{P}^\ell, \quad \text{for } \ell = 1, \dots, \ell_{\max} - 1. \tag{9}$$

Note that with this formula, computing the coarse versions of matrix $\mathbf{A}^1$ do not require constructing coarser meshes and performing the standard assembly operations at each one of the levels. In other words, the multigrid method is a purely algebraic procedure, once the transfer operators $\mathbf{P}^\ell$ and $\mathbf{R}^\ell$ are available.

A multigrid method is presented in Algorithm 1. The basic idea of the algorithm is to compute an approximation of the solution of the system of equations associated on the finest multigrid level $\ell = 1$ using the smoother $\boldsymbol{S}^1$, i.e. a cheap iterative method that is very efficient in reducing high frequency error components, and introducing an auxiliary system of equations on the next level $\ell = 2$, where remaining error components are again highly oscillatory, for computing a correction. This auxiliary system on level

$\ell = 2$ is approximated following the same idea by using the smoother $\boldsymbol{S}^2$ and introducing another auxiliary system at the next level $\ell = 3$. This coarsening procedure is repeated until the coarsest level $\ell = \ell_{\max}$ is reached. At the coarsest level, the problem can be efficiently solved with a direct method, and therefore, no further coarsening is needed. When the correction on the coarsest level is available, it is projected back through the different levels until the finest discretization is reached. The resulting method is called a V-cycle denoting the sequential order in which the different levels are visited, see Figure 1.

Procedure $\boldsymbol{\Phi}^{\mathrm{V-cycle}}(\mathbf{A}^\ell, \mathbf{b}^\ell, \mathbf{x}^\ell, \ell)$
- **Input:** System matrix $\mathbf{A}^\ell$, right hand side $\mathbf{b}^\ell$, initial guess $\mathbf{x}^\ell$, and current level ℓ.
- **Output:** Approximated solution of the system $\mathbf{A}^\ell \mathbf{x}^\ell = \mathbf{b}^\ell$.
- **if** $\ell < \ell_{\max}$ *(fine level)* **then**
  - Apply pre-smoother:
  - $\mathbf{x}^\ell \leftarrow \boldsymbol{S}^\ell(\mathbf{A}^\ell, \mathbf{b}^\ell, \mathbf{x}^\ell)$
  - Compute coarse level matrix:
  - $\mathbf{A}^{\ell+1} = \mathbf{R}^\ell \mathbf{A}^\ell \mathbf{P}^\ell$
  - Compute coarse level residual:
  - $\mathbf{b}^{\ell+1} = \mathbf{R}^\ell(\mathbf{b}^\ell - \mathbf{A}^\ell \mathbf{x}^\ell)$
  - Compute coarse level correction recursively:
  - $\mathbf{x}^{\ell+1} = \mathbf{0}$
  - $\mathbf{x}^{\ell+1} \leftarrow \boldsymbol{\Phi}^{\mathrm{V-cycle}}(\mathbf{A}^{\ell+1}, \mathbf{b}^{\ell+1}, \mathbf{x}^{\ell+1}, \ell+1)$
  - Apply correction:
  - $\mathbf{x}^\ell \leftarrow \mathbf{x}^\ell + \mathbf{P}^\ell \mathbf{x}^{\ell+1}$
  - Apply post-smoother:
  - $\mathbf{x}^\ell \leftarrow \boldsymbol{S}^\ell(\mathbf{A}^\ell, \mathbf{b}^\ell, \mathbf{x}^\ell)$
  - **return** $\mathbf{x}^\ell$
- **else** *(coarsest level)*
  - Apply a direct solver:
  - $\mathbf{x}^\ell \leftarrow \left(\mathbf{A}^\ell\right)^{-1} \mathbf{b}^\ell$
  - **return** $\mathbf{x}^\ell$

Algorithm 1: Multigrid V-cycle with $\ell_{\max}$ levels for the resolution of the fine level system $\mathbf{A}^1\mathbf{x}^1 = \mathbf{b}^1$.

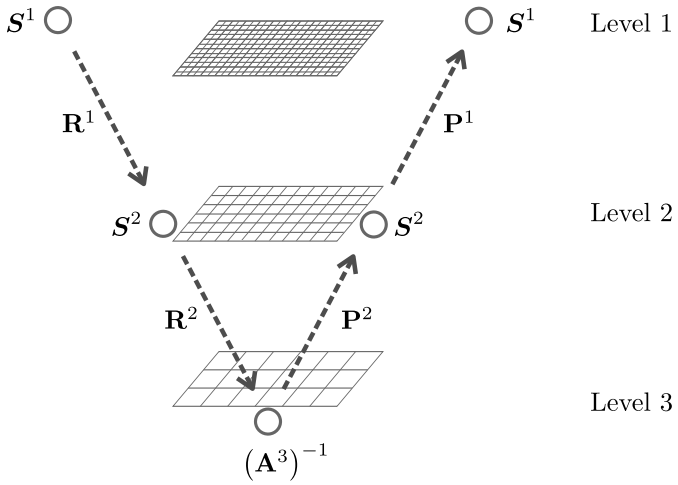

Figure 1: Illustration of a multigrid V-cycle with 3 levels.

Multigrid techniques are well established mainly for single-field problems. For elliptic problems (e.g. the steady-state heat equation or structural mechanics), multigrid methods are optimal in the sense that the cost of solving the system is proportional to the system size, cf. [43], and can be implemented in parallel rendering good scalability. The application of standard multigrid techniques (i.e. aggregation-

based methods [40]) to coupled problems is more challenging. Using standard aggregation-based AMG in coupled problems is known as *fully coupled* AMG and it is efficient for a number of applications including incompressible flows coupled with transport [44], resistive magneto hydrodynamics [13], and semiconductor device modeling [45]. However, the applicability of fully coupled AMG as a general purpose preconditioner for coupled problems is limited. This approach requires that the degrees of freedom of the underlying physical fields are collocated at the same mesh nodes, which is not fulfilled in many cases. That is, the method requires the same finite element interpolation for all field variables. Thus, more versatile methods are required for preconditioning coupled problems.

4 Unified framework for preconditing coupled problems

In the following, we present three general purpose preconditioners for solving the monolithic linear system (5). The presented methods are preconditioners presented previously in the literature with the common feature that they are able to be implemented for the generic coupled system (5). For the sake of simplicity, system (5) is written here as a generic system of equations with block structure:

$$\underbrace{\begin{bmatrix} \mathbf{A}_{11} & \cdots & \mathbf{A}_{1N} \\ \vdots & \ddots & \vdots \\ \mathbf{A}_{N1} & \cdots & \mathbf{A}_{NN} \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_N \end{bmatrix}}_{\mathbf{x}} = \underbrace{\begin{bmatrix} \mathbf{b}_1 \\ \vdots \\ \mathbf{b}_N \end{bmatrix}}_{\mathbf{b}}. \tag{10}$$

By setting $\mathbf{A}_{ij} := \partial \mathbf{f}_i / \partial \mathbf{u}_j$, $\mathbf{x}_i := \Delta \mathbf{u}_i$, and $\mathbf{b}_i := -\mathbf{f}_i$, equation (5) is recovered.

4.1 Block Gauss-Seidel (BGS) preconditioner

The first preconditioner is based on a standard BGS iteration for solving the system (10):

$$\mathbf{x}^{(k)} = \mathbf{x}^{(k-1)} + \mathbf{M}_{\mathrm{BGS}}^{-1}(\mathbf{b} - \mathbf{A}\mathbf{x}^{(k-1)}), \quad k = 1, \ldots, k_{\max}, \tag{11}$$

where $\mathbf{x}^{(k)}$ is the k-th iterate in the solution process, $\mathbf{x}^{(0)}$ is the initial guess and $k_{\max}$ is the selected number of iterations. The (forward) BGS iteration is recovered by setting $\mathbf{M}_{\mathrm{BGS}}$ as:

$$\mathbf{M}_{\mathrm{BGS}} := \begin{bmatrix} \mathbf{A}_{11} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{A}_{N1} & \cdots & \mathbf{A}_{NN} \end{bmatrix}.$$

Matrix $\mathbf{M}_{\mathrm{BGS}}$ is an approximation of matrix $\mathbf{A}$ obtained by discarding the blocks placed at the upper triangle. The same idea applies by discarding the lower triangle instead of the upper (referred as backward BGS) or alternating the lower and upper triangles at each iteration (referred to as symmetric BGS). In our case, the three versions are implemented so that the user can choose the best variant for the particular application. The quality of the BGS preconditioner strongly depends in the error committed when discarding the upper (or lower) off-diagonal blocks of matrix $\mathbf{A}$. Thus, the BGS method requires that the diagonal block are dominant, see e.g. [46].

The forward BGS iteration (11) is written in more detail as

$$\mathbf{x}_i^{(k)} = \mathbf{A}_{ii}^{-1} \left[\mathbf{b}_i - \sum_{j=1}^{i-1} \mathbf{A}_{ij} \mathbf{x}_j^{(k)} + \sum_{j=i+1}^{N} \mathbf{A}_{ij} \mathbf{x}_j^{(k-1)} \right] \quad \text{for } i = 1, \ldots, N. \tag{12}$$

Thanks to the triangular structure of $\mathbf{M}_{\mathrm{BGS}}$, computing the iteration (12) only requires solving systems with the main diagonal blocks $\mathbf{A}_{ii}$, $i = 1, \ldots, N$. That is, the BGS iteration transforms the coupled

problem (10) into a set of uncoupled single field-problems. The advantage of this approach is that the resulting uncoupled systems can be solved with standard methods available for single field-problems, e.g. AMG methods tailored to each of the underlying fields. If AMG is considered for solving the uncoupled problems, the resulting method is denoted as BGS(AMG), indicating that the preconditioner is constructed using two main ingredients: an outer BGS iteration plus inner uncoupled AMG methods (see Figure 2). In our implementation, we allow the user to chose from several different methods for solving the single-field problems. The available methods include aggregation-based AMG (provided by the MueLu package [41]), relaxation methods and incomplete factorizations (both provided by the Ifpack package [47]), and direct solvers (provided by the Amesos package [48]).

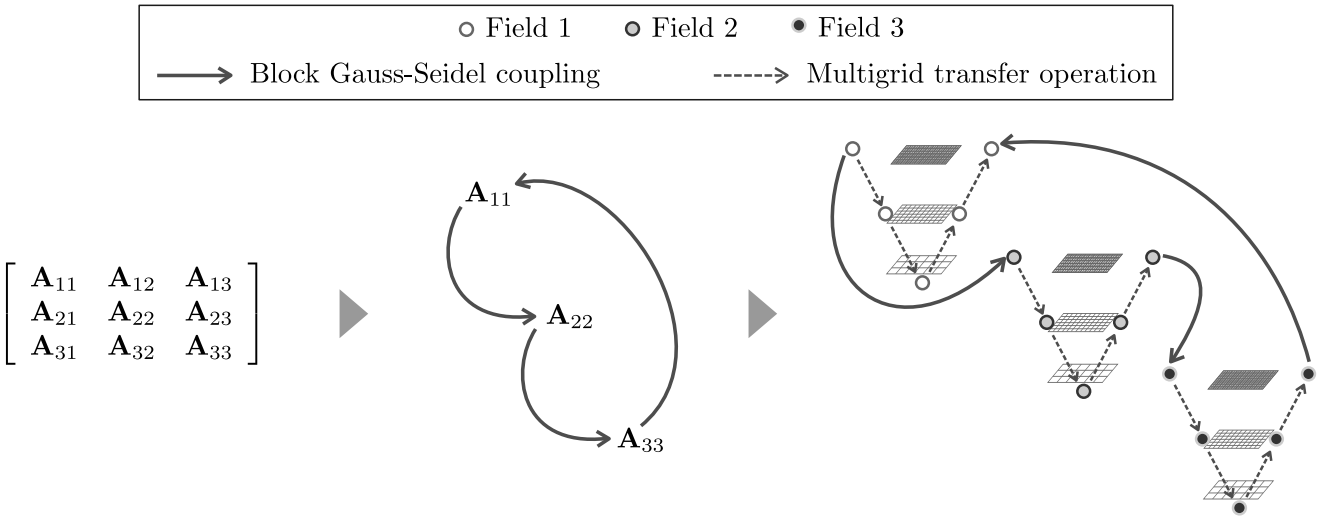

Figure 2: Illustration of the BGS(AMG) preconditioner for a coupled problem with 3 fields. The method is built using to main steps: First a Block Gauss-Seidel (BGS) iteration is considered for handling the coupling. Second, the resulting uncoupled systems are attacked with standard algebraic multigrid (AMG) methods.

The BGS(AMG) is a general solver in the sense that it can be implemented for an arbitrary number of fields, and therefore, it can be potentially applied to a wide range of coupled problems. However, the method requires that all the main diagonal matrices $\mathbf{A}_{ii}$ are non-singular, which is not always fulfilled (e.g. in saddle point problems). Another potential drawback for some problems is that coupling is only achieved at the finest level, as it can easily be seen from Figure 2.

4.2 SIMPLE preconditioner for an arbitrary number of fields

The next preconditioner is based on an idea similar to the semi-implicit method for pressure-linked equations (SIMPLE) [19] which was originally proposed for solving the saddle point structure of the Navier-Stokes equations. This method is required in situations where the BGS(AMG) preconditioner fails (e.g. when one of the matrices $\mathbf{A}_{ii}$ is singular or zero). The SIMPLE method is designed for problems with a 2×2 block structure like:

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{bmatrix}. \tag{13}$$

The nice property of the method is that the main diagonal matrix $\mathbf{A}_{22}$ can be singular or even zero, as it turns out in saddle point problems, see [18].

The SIMPLE method for solving system (13) is the solver characterized by the iteration

$$\begin{bmatrix} \mathbf{x}_1^{(k)} \\ \mathbf{x}_2^{(k)} \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1^{(k-1)} \\ \mathbf{x}_2^{(k-1)} \end{bmatrix} + \mathbf{M}_{\text{SIMPLE}}^{-1} \left(\begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{bmatrix} - \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{x}_1^{(k-1)} \\ \mathbf{x}_2^{(k-1)} \end{bmatrix} \right), \quad k = 1, \dots, k_{\max}, \tag{14}$$

where the superscript $[\cdot]^{(k)}$ denotes the k-th iteration in the solution process and the matrix $\mathbf{M}_{\text{SIMPLE}}$ is the following block factorization of the system matrix $\mathbf{A}$:

$$\mathbf{M}_{\text{SIMPLE}} := \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} \\ \mathbf{A}_{21} & \widetilde{\mathbf{S}} \end{bmatrix} \begin{bmatrix} \mathbf{I} & (\widetilde{\mathbf{A}}_{11})^{-1}\mathbf{A}_{12} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}. \tag{15}$$

Here, matrix $\widetilde{\mathbf{A}}_{11}$ is an easy-to-invert approximation of $\mathbf{A}_{11}$ and $\widetilde{\mathbf{S}} = \mathbf{A}_{22} - \mathbf{A}_{21}(\widetilde{\mathbf{A}}_{11})^{-1}\mathbf{A}_{12}$ is the associated approximation of the Schur complement matrix. If $\widetilde{\mathbf{A}}_{11} = \mathbf{A}_{11}$, matrix $\mathbf{M}_{\text{SIMPLE}}$ coincides with the system matrix $\mathbf{A}$ and the method (14) converges in one iteration. Thus, the effectivity of the SIMPLE iteration depends on the quality of the approximation $\widetilde{\mathbf{A}}_{11}$. In the numerical examples, we construct $\widetilde{\mathbf{A}}_{11}$ following the so-called SIMPLEC [49] approach which consists in taking $\widetilde{\mathbf{A}}_{11}$ as the diagonal matrix whose entries are the absolute values of the row sums of $\mathbf{A}_{11}$.

Iteration (14) is written in more detail as

$$\Delta\mathbf{x}_1^{(k-1)} = \mathbf{A}_{11}^{-1}\left[\mathbf{b}_1 - \mathbf{A}_{12}\mathbf{x}_2^{(k-1)}\right], \quad \text{(Predictor equation)} \tag{16a}$$

$$\Delta\mathbf{x}_2^{(k-1)} = \widetilde{\mathbf{S}}^{-1}\left[\mathbf{b}_2 - \mathbf{A}_{21}\Delta\mathbf{x}_1^{(k-1)} - \mathbf{A}_{22}\mathbf{x}_2^{(k-1)}\right], \quad \text{(Schur equation)} \tag{16b}$$

$$\mathbf{x}_2^{(k)} = \mathbf{x}_2^{(k-1)} + \Delta\mathbf{x}_2^{(k-1)}, \quad \text{(Update)} \tag{16c}$$

$$\mathbf{x}_1^{(k)} = \mathbf{x}_1^{(k-1)} - (\widetilde{\mathbf{A}}_{11})^{-1}\mathbf{A}_{12}\Delta\mathbf{x}_2^{(k-1)}. \quad \text{(Update)} \tag{16d}$$

Steps (16a) and (16b) are known as the predictor and Schur equations and require solving systems with the matrices $\mathbf{A}_{11}$ and $\widetilde{\mathbf{S}}$ respectively. Thus, the SIMPLE method can be seen as a procedure that transforms a coupled problem (13) into a set of two uncoupled systems that can be solved with standard techniques for single-field problems, e.g. AMG methods. The resulting preconditioner is denoted here as SIMPLE(AMG) highlighting that the methods consists in an outer SIMPLE iteration for uncoupling the fields and inner AMG solvers for the resulting uncoupled systems, see Figure 3.

The SIMPLE method can be extended to cope with generic coupled problems like (10) with more than two physical fields (i.e. $N > 2$). To this end, we follow a similar strategy to the one proposed by Badia et al. [5] in the context of magneto-hydro dynamics. The methodology consists in merging the blocks of the system (10) in order to recover a 2×2 block structure. For instance, a coupled problem with a 3×3 block structure can be rearranged into a 2×2 structure by merging the first and second fields:

$$\left[\begin{array}{cc|c} \mathbf{A}_{11} & \mathbf{A}_{12} & \mathbf{A}_{13} \\ \mathbf{A}_{21} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \hline \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \end{array}\right] \left[\begin{array}{c} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \hline \mathbf{x}_3 \end{array}\right] = \left[\begin{array}{c} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \hline \mathbf{b}_3 \end{array}\right] \quad \longrightarrow \quad \begin{bmatrix} \overline{\mathbf{A}}_{11} & \overline{\mathbf{A}}_{12} \\ \overline{\mathbf{A}}_{21} & \overline{\mathbf{A}}_{22} \end{bmatrix} \begin{bmatrix} \overline{\mathbf{x}}_1 \\ \overline{\mathbf{x}}_2 \end{bmatrix} = \begin{bmatrix} \overline{\mathbf{b}}_1 \\ \overline{\mathbf{b}}_2 \end{bmatrix}, \tag{17}$$

where

$$\overline{\mathbf{A}}_{11} := \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}, \quad \overline{\mathbf{A}}_{12} := \begin{bmatrix} \mathbf{A}_{13} \\ \mathbf{A}_{23} \end{bmatrix}, \quad \overline{\mathbf{A}}_{21} := \begin{bmatrix} \mathbf{A}_{31} & \mathbf{A}_{32} \end{bmatrix}, \quad \overline{\mathbf{A}}_{22} := \mathbf{A}_{33},$$

and

$$\overline{\mathbf{x}}_1 := \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix}, \quad \overline{\mathbf{x}}_2 := \mathbf{x}_3, \quad \overline{\mathbf{b}}_1 := \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{bmatrix}, \quad \overline{\mathbf{b}}_2 := \mathbf{b}_3.$$

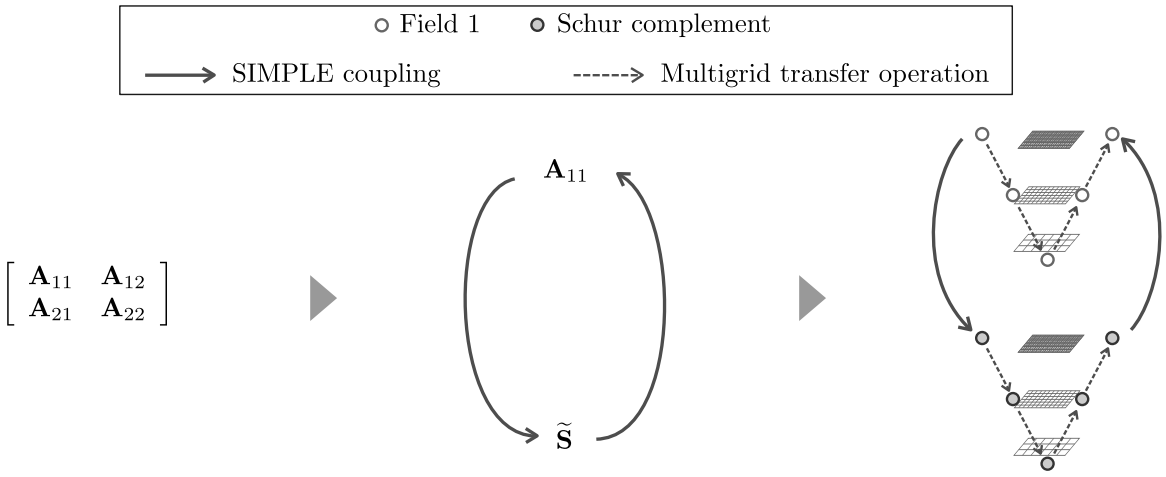

Figure 3: Illustration of the SIMPLE(AMG) preconditioner for a coupled system with 2 fields. The method is obtained by applying an outer SIMPLE iteration for uncoupling the fields and AMG methods for solving the resulting uncoupled single-field systems.

The merging procedure has to be performed carefully, because after applying the SIMPLE iteration (16), the resulting predictor and Schur equations (16a) and (16b) might involve more than one physical fields, and therefore standard AMG cannot be applied directly. For example, the merging procedure for the 3×3 problem given in equation (17) leads to a system matrix for the predictor equation, which involves two physical fields:

$$\overline{\mathbf{A}}_{11} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}.$$

Thus, in order to solve the system involving this matrix we cannot directly use standard AMG. Instead, we consider a block iterative method, either BGS or SIMPLE, for uncoupling the fields. This can be seen as a recursive application of the different proposed approaches, which results in uncoupled systems that can be solved with standard AMG methods, see Figure 4.

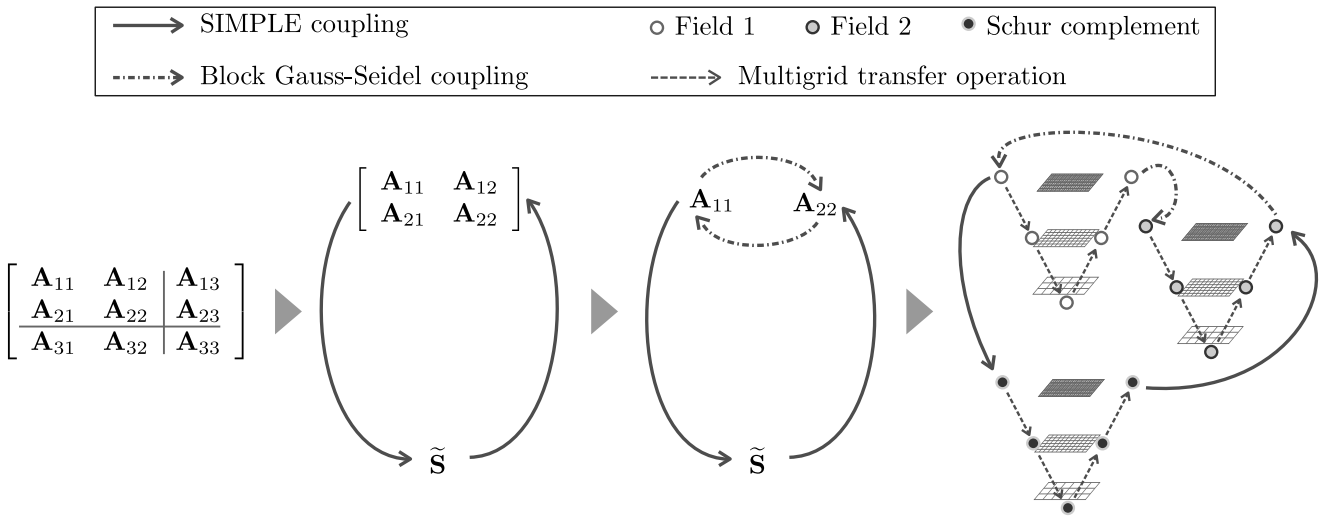

Figure 4: Using the SIMPLE method in problems with more than 2 physical fields. The blocks of the original system are merged in order to recover the 2×2 block structure required by the SIMPLE method. This results in a coupled problem to be solved within each SIMPLE iteration. A block iteration (e.g. BGS) is considered for uncoupling this problem, which results in a set of uncoupled single-field problems that can be efficiently solved with AMG.

4.3 Monolithic AMG preconditioner

The third general purpose method presented here is based on the so-called monolithic AMG preconditioner proposed by Gee et al. [8] for fluid-structure interaction (FSI) applications (see also Mayr et al. [50]). Note that both the BGS(AMG) and SIMPLE(AMG) methods described above transform the coupled problem in a set of uncoupled systems which are solved with independent AMG methods. The major drawback of this approach is that the coupling between the fields is only resolved on the finest multigrid level through the outer block iteration. Even if efficient inner AMG solvers are used, the coupling can be under-resolved if not enough iterations of the outer BGS or SIMPLE are carried out. This drawback is overcome by Gee et al. [8] in the context of FSI with a multigrid method which accounts for the coupling at all multigrid levels, while keeping the block structure of the fine level matrix. The monolithic AMG method is extended here for a generic coupled problem like (10).

The main ingredients of the monolithic AMG method are standard AMG solvers for each one of the underlying physical fields. Let us assume that we can build specific AMG methods for each one of the fields using information of the matrices $\mathbf{A}_{ii}$, $i = 1, \dots, N$. From the computed AMG hierarchies, we extract the usual multigrid objects: 1) the projection operators $\mathbf{P}_i^\ell$, 2) the restriction operators $\mathbf{R}_i^\ell$, and 3) the smoothers $\boldsymbol{S}_i^\ell$. The subscript $[\cdot]_i$ stands for each one of the fields $i = 1, \dots, N$, and the superscript $[\cdot]^\ell$ stands for each one of the levels $\ell = 1, \dots, \ell_{\max}$. For the sake of simplicity, we assume that all the number of levels $\ell_{\max}$ is the same across all the fields.

The key idea of monolithic AMG is to reuse these single-field multigrid objects to build a multigrid method for the entire coupled problem. The transfer operators $\mathbf{P}^\ell$ and $\mathbf{R}^\ell$ for the monolithic AMG method are defined as the block diagonal matrices

$$\mathbf{P}^\ell := \begin{bmatrix} \mathbf{P}_1^\ell & & \\ & \ddots & \\ & & \mathbf{P}_N^\ell \end{bmatrix}, \quad \mathbf{R}^\ell := \begin{bmatrix} \mathbf{R}_1^\ell & & \\ & \ddots & \\ & & \mathbf{R}_N^\ell \end{bmatrix}, \quad \ell = 1, \dots, \ell_{\max} - 1, \tag{18}$$

which can be computed from the transfer operators $\mathbf{P}_i^\ell$ and $\mathbf{R}_i^\ell$ previously computed for the underlying single-field problems.

The transfer operators $\mathbf{P}^\ell$ and $\mathbf{R}^\ell$ for the entire coupled problem have the nice property that lead to coarser system matrices $\mathbf{A}^\ell$ having the same block structure as the fine level matrix $\mathbf{A}^1$. The coarse level matrices are obtained as

$$\underbrace{\begin{bmatrix} \mathbf{A}_{11}^{\ell+1} & \cdots & \mathbf{A}_{1N}^{\ell+1} \\ \vdots & \ddots & \vdots \\ \mathbf{A}_{N1}^{\ell+1} & \cdots & \mathbf{A}_{NN}^{\ell+1} \end{bmatrix}}_{\mathbf{A}^{\ell+1}} := \underbrace{\begin{bmatrix} \mathbf{R}_1^\ell & & \\ & \ddots & \\ & & \mathbf{R}_N^\ell \end{bmatrix}}_{\mathbf{R}^\ell} \underbrace{\begin{bmatrix} \mathbf{A}_{11}^\ell & \cdots & \mathbf{A}_{1N}^\ell \\ \vdots & \ddots & \vdots \\ \mathbf{A}_{N1}^\ell & \cdots & \mathbf{A}_{NN}^\ell \end{bmatrix}}_{\mathbf{A}^\ell} \underbrace{\begin{bmatrix} \mathbf{P}_1^\ell & & \\ & \ddots & \\ & & \mathbf{P}_N^\ell \end{bmatrix}}_{\mathbf{P}^\ell}, \tag{19}$$

$\ell = 1, \dots, \ell_{\max} - 1$, starting from the fine level matrix $\mathbf{A}^1$. Thanks to the diagonal structure of operators $\mathbf{P}^\ell$ and $\mathbf{R}^\ell$, the blocks of the coarse level matrices $\mathbf{A}^{\ell+1}$ are computed form the blocks of the fine level matrix $\mathbf{A}^\ell$, namely $\mathbf{A}_{ij}^{\ell+1} = \mathbf{R}_i^\ell \mathbf{A}_{ij}^\ell \mathbf{P}_j^\ell$, without mixing up the fields. That is, the physical interpretation of the coarse level blocks $\mathbf{A}_{ij}^{\ell+1}$ is the same as the fine level blocks $\mathbf{A}_{ij}^\ell$. This implies that all coarse level systems with matrices $\mathbf{A}^\ell$ have the same block structure as the fine level problem (10), see Figure 5.

Finally, the smoothers $\mathbf{S}^\ell$ used within the monolithic AMG method have to be defined. Since the coarse level matrices retain the block structure of the fine level problem, a block iterative method (i.e. BGS or SIMPLE) is used. More precisely, the smoothers $\boldsymbol{S}^\ell$ are build as an outer iteration (i.e. BGS or SIMPLE) for uncoupling the blocks in the system matrices $\mathbf{A}^\ell$, and the resulting single-field uncoupled problems are solved with the smothers $\boldsymbol{S}_i^\ell$ previously extracted from the single-field multigrid hierarchies, see Figure 5.

The monolithic AMG method is finally constructed as a standard V-cycle using the projectors $\mathbf{P}^\ell$, restrictors $\mathbf{R}^\ell$, coarse level matrices $\mathbf{A}^\ell$, and smoothers $\mathbf{S}^\ell$, see Figure 5. This method is referred to as

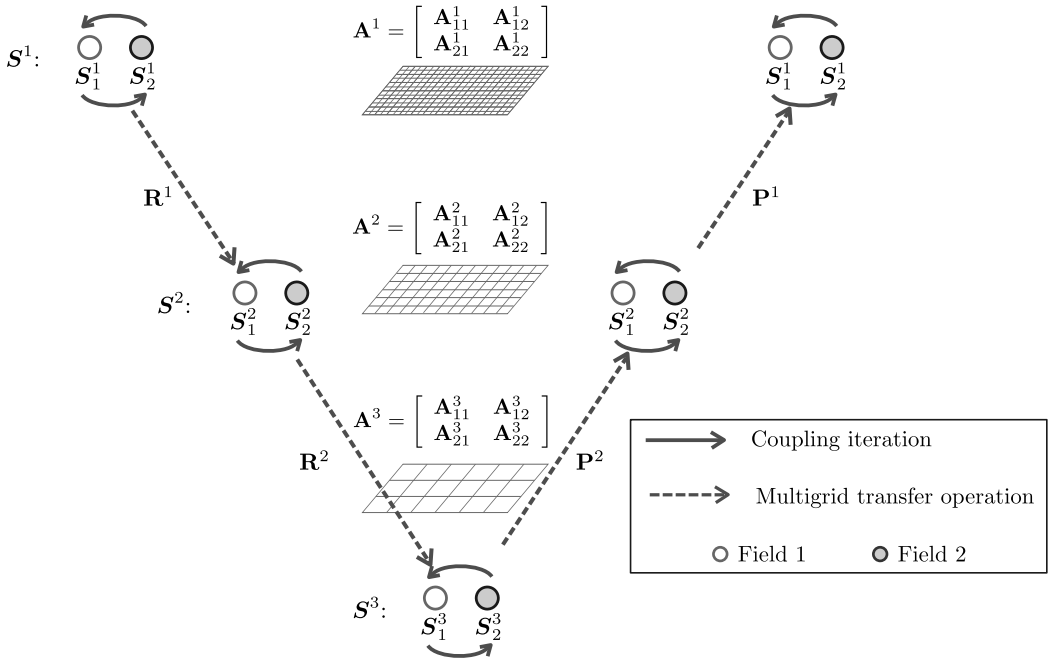

Figure 5: Illustration of a monolithic AMG method for a coupled problem with 2 fields.

AMG(BGS) or AMG(SIMPLE) denoting that the block iteration (BGS or SIMPLE) is performed within each level of the AMG hierarchy.

5 Implementation remarks

The presented preconditioners for coupled problems (i.e. BGS, SIMPLE, and monolithic AMG) are implemented in such a way that they can be combined in order to build other more complex preconditioners allowing to solve challenging applications. For instance, any of these three methods can be used within the SIMPLE iteration (16) to solve the predictor equation (16a) or the Schur complement equation (16a) if the problems to be solved involve more than one physical field. This can be seen as a recursive application of a block preconditioner within another block preconditioner requiring a flexible implementation. We benefit from this flexibility later in Section 6.3 for building preconditioners for the lung model, where we consider both the BGS method and the monolithic AMG preconditioner within a SIMPLE iteration.

The flexible implementation is mainly achieved considering an object-oriented software design, see Figure 6. The framework is implemented using an interface class from which the BGS, SIMPLE, and monolithic AMG methods are derived. The solvers used within the SIMPLE iteration (16) to solve the predictor equation (16a) or the Schur complement equation (16a) are declared using the interface class, which allows calling BGS, SIMPLE, or monolithic AMG methods within the SIMPLE iteration. Similarly, the level smoothers within the monolithic AMG method are also declared using an interface class, allowing to have BGS or SIMPLE as smoothers. From the user point of view, the preconditioners are selected using a high level interface without the need of being familiar with the technical details of the underlying implementation.

6 Applications

In this section, we describe the different application examples that we have chosen to demonstrate the applicability and quality of the proposed approach.

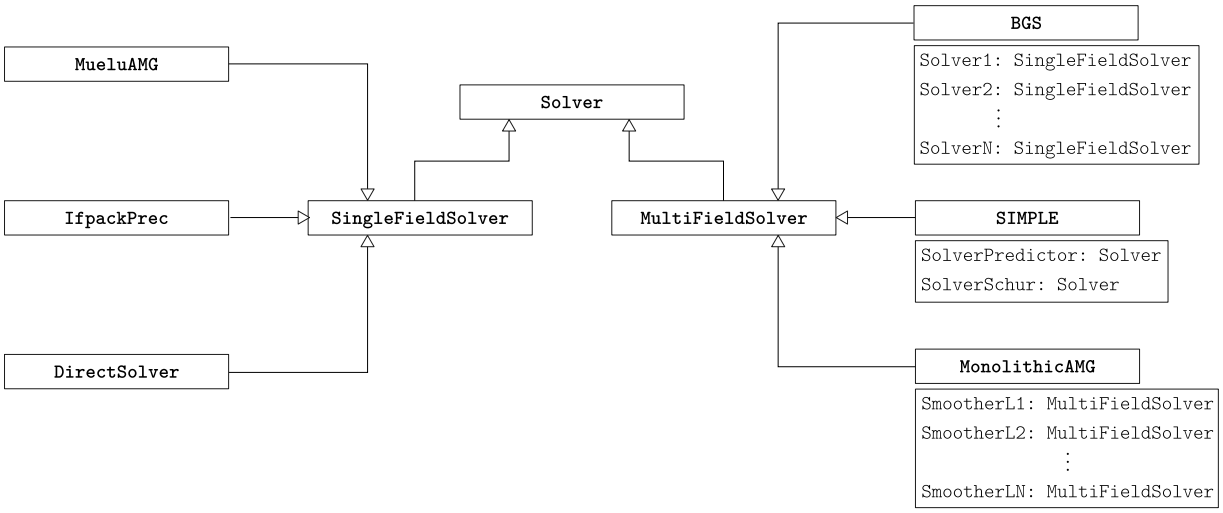

Figure 6: Simplified class diagram of the object oriented design of the proposed framework.

6.1 Thermo-structure interaction (TSI)

The first problem type in which the proposed methods are tested is the interaction between an elastic solid and a thermal field. For the sake of a simple presentation, the solid is modeled here as a linear elastic material under small deformations, but the same approach holds for more sophisticated structural models. The solid is governed by the equations of elastodynamics

$$\rho\ddot{\boldsymbol{d}} - \nabla\cdot\boldsymbol{\sigma}(\boldsymbol{d},u) = \boldsymbol{b} \quad \Omega\times I, \tag{20}$$

describing the mechanical equilibrium in the domain Ω over the time interval I. The unknown displacements and accelerations are $\boldsymbol{d}$ and $\ddot{\boldsymbol{d}}$ respectively, $\boldsymbol{\sigma}$ stands for the unknown stresses, ρ is the density, and $\boldsymbol{b}$ is an arbitrary body force.

Equation (20) is augmented with the constitutive model

$$\boldsymbol{\sigma}(\boldsymbol{d},u) := \boldsymbol{\mathcal{C}} : \boldsymbol{\varepsilon}(\boldsymbol{d}) + m(u-u_0)\boldsymbol{I}, \tag{21}$$

where the term $\boldsymbol{\mathcal{C}} : \boldsymbol{\varepsilon}(\boldsymbol{d})$ is the standard elastic stress associated with the strains $\boldsymbol{\varepsilon}(\boldsymbol{d})$ via the fourth order elasticity tensor $\boldsymbol{\mathcal{C}}$, and the term $m(u-u_0)\boldsymbol{I}$ is the thermal stress introducing the influence of the temperature field. The value u_0 is the initial temperature, $\boldsymbol{I}$ is the second order identity tensor, and m is defined as $m := -(1\mu+3\lambda)\alpha$, where λ and μ are the Lamé constants, and α is the so-called *coefficient of thermal expansion*.

The TSI problem is obtained by coupling the mechanical equations (20) and (21) with the time-dependent heat equation:

$$\rho C\dot{u} - \nabla\cdot\boldsymbol{q}(u) - r(\dot{\boldsymbol{d}})u = s \quad \Omega\times I, \tag{22}$$

where C is the heat capacity, $\boldsymbol{q}(u) = -\mathbf{k}\nabla u$ is the heat flux, $\mathbf{k}$ is the second order tensor describing the thermal conductivity, and s is an arbitrary source term. The influence of the mechanical displacements in the heat equation is modeled by the deformation-dependent reaction coefficient

$$r(\dot{\boldsymbol{d}}) := m\left(\boldsymbol{I} : \boldsymbol{\varepsilon}(\dot{\boldsymbol{d}})\right).$$

For the sake of brevity, it is assumed that the structural equation (20) and the heat equation (22) are closed with the usual boundary and initial conditions.

Equations (20) and (22) lead to a fully coupled problem: the structural equation depends on the temperature and vice versa. Several solution strategies are proposed in the literature for resolving this coupling.

They include staggered [51, 52, 53, 54] and monolithic methods [7]. In strongly coupled settings (i.e. under strong values of the thermal expansion coefficient α), the work by Danowski et al. [7] shows that monolithic methods are often superior to partitioned schemes.

The monolithic scheme introduced by Danowski et al. [7] leads to a system of equations with block structure requiring special preconditioners as the ones previously presented. The derivation of this linear system of equations is briefly presented here, see [7] for further details.

The discretization in space and time of equations (20) and (22) lead to a set of non-linear equations that have to be solved at each time step in order to obtain the displacements and temperatures. These equations are denoted as

$$\mathbf{f}_{\mathrm{S}}^{n}(\mathbf{d}^{n}, \mathbf{u}^{n}) = \mathbf{0}, \tag{23a}$$

$$\mathbf{f}_{\mathrm{T}}^{n}(\mathbf{d}^{n}, \mathbf{u}^{n}) = \mathbf{0}, \tag{23b}$$

where $\mathbf{f}_{\mathrm{S}}^{n}$ and $\mathbf{f}_{\mathrm{T}}^{n}$ are the discrete structural and thermal equations, and $\mathbf{d}^{n}$ and $\mathbf{u}^{n}$ are the discrete displacements and temperatures at a generic time point t_n. Equations (23) are a discrete, fully coupled, non-linear problem which can be written more compactly as

$$\mathbf{f}^{\mathrm{TSI}}(\mathbf{x}^{\mathrm{TSI}}) = \mathbf{0}, \quad \text{with} \quad \mathbf{f}^{\mathrm{TSI}} = \begin{bmatrix} \mathbf{f}_{\mathrm{S}} \\ \mathbf{f}_{\mathrm{T}} \end{bmatrix} \text{ and } \mathbf{x}^{\mathrm{TSI}} = \begin{bmatrix} \mathbf{d} \\ \mathbf{u} \end{bmatrix},$$

where the superscript $[\cdot]^n$ denoting the time point is dropped for simplicity. Following the monolithic approach by Danowski et al. [7], this non-linear problem is solved using a Newton-type method

$$\mathbf{x}_{(i+1)}^{\mathrm{TSI}} = \mathbf{x}_{(i)}^{\mathrm{TSI}} + \Delta\mathbf{x}_{(i)}^{\mathrm{TSI}} \quad \text{with} \quad \left[\frac{\partial \mathbf{f}^{\mathrm{TSI}}}{\partial \mathbf{x}^{\mathrm{TSI}}}\right]_{(i)} \Delta\mathbf{x}_{(i)}^{\mathrm{TSI}} = -\mathbf{f}_{(i)}^{\mathrm{TSI}}, \tag{24}$$

where the subscript $[\cdot]_{(i)}$ denotes the Newton iteration. The system of equations in (24) involves the linearization of the TSI equations ($\mathbf{f}_{\mathrm{S}}$ and $\mathbf{f}_{\mathrm{T}}$) with respect the TSI variables ($\mathbf{d}$ and $\mathbf{u}$). Thus, the system of linear equations to be solved at each Newton step has a 2×2 block structure associated with the solid and thermal fields:

$$\begin{bmatrix} \dfrac{\partial \mathbf{f}_{\mathrm{S}}}{\partial \mathbf{d}} & \dfrac{\partial \mathbf{f}_{\mathrm{S}}}{\partial \mathbf{u}} \\ \dfrac{\partial \mathbf{f}_{\mathrm{T}}}{\partial \mathbf{d}} & \dfrac{\partial \mathbf{f}_{\mathrm{T}}}{\partial \mathbf{u}} \end{bmatrix}_{(i)} \begin{bmatrix} \Delta\mathbf{d} \\ \Delta\mathbf{u} \end{bmatrix}_{(i)} = - \begin{bmatrix} \mathbf{f}_{\mathrm{S}} \\ \mathbf{f}_{\mathrm{T}} \end{bmatrix}_{(i)}. \tag{25}$$

The main diagonal blocks of the system are the usual matrices arising in structural mechanics and in the heat equation, and therefore, they are available in many finite element packages. The off-diagonal blocks are associated with the coupling terms in equations (20) and (22) and are computed as indicated in [7].

Name	Description of the preconditioner
BGS(AMG)	Preconditioner described in Section 4.1, using a backward BGS iteration.
SIMPLE(AMG)	Preconditioner described in Section 4.2, taking the structural equations as the predictor field and the thermal equations as Schur field.
AMG(BGS)	Preconditioner described in Section 4.3 with backward BGS smoothers.

Table 1: Description of the preconditioners for the monolithic thermo-structure interaction problem.

The monolithic system (25) requires often to be preconditioned in order to be solved with a Krylov method efficiently. Since this system is a particular case of the general coupled problem (10), we can use the generic preconditioners previously presented, see Table 1. The first method in Table 1 is the BGS(AMG) preconditioner presented in Section 4.1 using a backwards BGS iteration. We choose a

backward BGS since the effect of the structural displacement on the thermal equations is less important than the opposite effect, see [7]. The second method is the SIMPLE(AMG) preconditioner introduced in Section 4.2. Finally, the third preconditioner is the monolithic AMG method presented in Section 4.3. We consider a backwards BGS method as smoother within the AMG hierarchy and for that reason the method is denoted as AMG(BGS).

Among the three preconditioners given in Table 1, the method BGS(AMG) corresponds to the one considered in the work by Danowski et al. [7]. To the knowledge of the authors, the other two techniques are not previously used in the context of TSI.

6.2 Fluid-structure interaction (FSI)

The second coupled problem considered in this work is FSI. We adopt the monolithic approach proposed by Küttler and Wall [55] which furnishes challenging linear systems with block structure. For the sake of completeness, we briefly summarize the method here.

Let us consider the coupling between a solid and a fluid at the interface Γ. The displacements $\boldsymbol{d}^{\mathrm{S}}$ of the solid are governed by the equations of structural mechanics in the solid domain Ω^{S}. On the other hand, the fluid velocities $\boldsymbol{u}^{\mathrm{F}}$ and pressure p are governed by the arbitrary Lagrangian-Eulerian (ALE) version of the Naiver-Stokes equations [56] in the fluid domain Ω^{F}. The ALE approach requires introducing a third (non-physical) grid movement field $\boldsymbol{d}^{\mathrm{G}}$ in Ω^{F} in order to accommodate the fluid mesh to deformations of the FSI interface. The mesh movement $\boldsymbol{d}^{\mathrm{G}}$ is an arbitrary but unique extension of the solid displacement into the fluid domain Ω^{F}.

After discretizing in space and time, the FSI problem reduces to find the discrete structural displacements $\mathbf{d}_n^{\mathrm{S}}$ in Ω^{S}, the discrete grid displacements $\mathbf{d}_n^{\mathrm{G}}$ in Ω^{F}, and the discrete fluid velocities $\mathbf{u}_n^{\mathrm{F}}$ and discrete fluid pressures $\mathbf{p}_n^{\mathrm{F}}$ in Ω^{F}, for each time point t_n.

These unknowns are determined by solving the non-linear equations resulting from the discretization of the structural, fluid and grid motion fields. These equations are denoted here as:

$$\mathbf{f}_{\mathrm{I},n}^{\mathrm{S}}(\mathbf{d}_n^{\mathrm{S}}) = \mathbf{0}, \tag{26a}$$

$$\mathbf{f}_{\mathrm{I},n}^{\mathrm{G}}(\mathbf{d}_n^{\mathrm{G}}) = \mathbf{0}, \tag{26b}$$

$$\mathbf{f}_{\mathrm{I},n}^{\mathrm{F}}(\mathbf{u}_n^{\mathrm{F}}, \mathbf{p}_n^{\mathrm{F}}, \mathbf{d}_n^{\mathrm{G}}) = \mathbf{0}, \tag{26c}$$

where the superscripts $[\cdot]^{\mathrm{S}}$, $[\cdot]^{\mathrm{G}}$ and $[\cdot]^{\mathrm{F}}$ stand for the structural field, the grid motion and the fluid respectively. The subscript $[\cdot]_{\mathrm{I}}$ denotes that the equations are associated only with mesh nodes in the interior of the respective domains (i.e. not at the FSI interface).

Equations (26) have to be augmented with boundary and initial conditions together with the FSI conditions at the interface Γ in order to have a well posed problem. Boundary and initial conditions are standard, and therefore, they are not discussed here. The coupling conditions at the FSI interface Γ are:

$$\boldsymbol{u}^{\mathrm{F}} = \dot{\boldsymbol{d}}^{\mathrm{S}} \quad \text{on } \Gamma \times I, \tag{27a}$$

$$\boldsymbol{d}^{\mathrm{G}} = \boldsymbol{d}^{\mathrm{S}} \quad \text{on } \Gamma \times I, \tag{27b}$$

$$\boldsymbol{\sigma}^{\mathrm{S}} \cdot \boldsymbol{n} = \boldsymbol{\sigma}^{\mathrm{F}} \cdot \boldsymbol{n} \quad \text{on } \Gamma \times I. \tag{27c}$$

The first condition is the kinematic continuity between fluid and structural velocities, the second is the continuity between structural displacements and the grid motion field, and the third is the equilibrium of forces at the FSI interface. The tensors $\boldsymbol{\sigma}^{\mathrm{S}}$ and $\boldsymbol{\sigma}^{\mathrm{F}}$ stand for the Cauchy stress of the solid and the fluid respectively, and $\boldsymbol{n}$ is a unique normal defined on the FSI interface.

Defining the discrete FSI problem requires introducing the discrete version of the FSI coupling condition (27c) which is denoted here as

$$\mathbf{f}_{\Gamma,n}^{\mathrm{S}}(\mathbf{d}_n^{\mathrm{S}}) + \mathbf{f}_{\Gamma,n}^{\mathrm{F}}(\mathbf{u}_n^{\mathrm{F}}, \mathbf{p}_n^{\mathrm{F}}, \mathbf{d}_n^{\mathrm{G}}) = \mathbf{0}, \tag{28}$$

where the subscript $[\cdot]_\Gamma$ denotes quantities associated with the FSI interface Γ. The vector $\mathbf{f}^{\mathrm{S}}_{\Gamma,n}$ represents the structural interface forces obtained by weighting (the discrete version of) $\boldsymbol{\sigma}^{\mathrm{S}} \cdot \boldsymbol{n}$ by the shape functions associated with the nodes on Γ. The same applies for the fluid interface forces $\mathbf{f}^{\mathrm{F}}_{\Gamma,n}$.

The interior field equations (26) together with the interface condition (28) form the fully coupled FSI problem to be solved at each time step:

$$\mathbf{f}^{\mathrm{FSI}}(\mathbf{x}^{\mathrm{FSI}}) = \mathbf{0} \quad \text{with } \mathbf{f}^{\mathrm{FSI}} := \begin{bmatrix} \mathbf{f}^{\mathrm{S}}_{\mathrm{I}} \\ \mathbf{f}^{\mathrm{G}}_{\mathrm{I}} \\ \mathbf{f}^{\mathrm{F}}_{\mathrm{I}} \\ \mathbf{f}^{\mathrm{S}}_{\Gamma} + \mathbf{f}^{\mathrm{F}}_{\Gamma} \end{bmatrix} \text{ and } \mathbf{x}^{\mathrm{FSI}} := \begin{bmatrix} \mathbf{d}^{\mathrm{S}}_{\mathrm{I}} \\ \mathbf{d}^{\mathrm{G}}_{\mathrm{I}} \\ \mathbf{u}^{\mathrm{F}}_{\mathrm{I}} \\ \mathbf{u}^{\mathrm{F}}_{\Gamma} \end{bmatrix}, \tag{29}$$

where the subscript $[\cdot]_n$ for the time step has been dropped for simplicity. Note that the FSI variable $\mathbf{x}^{\mathrm{FSI}}$ does not include the displacements $\mathbf{d}^{\mathrm{S}}_{\Gamma,n}$ nor $\mathbf{d}^{\mathrm{G}}_{\Gamma,n}$. This is because these displacements can be computed once $\mathbf{u}^{\mathrm{F}}_{\Gamma,n}$ is available using the coupling conditions (27a) and (27b). For the ease of notation, the fluid pressure $\mathbf{p}^{\mathrm{F}}$ is also dropped from the formula and it is assumed that the fluid quantities $\mathbf{u}^{\mathrm{F}}_{\mathrm{I}}$ and $\mathbf{u}^{\mathrm{F}}_{\Gamma}$ include the pressure degrees of freedom.

By merging the interior and interface fluid degrees of freedom, namely

$$\mathbf{u}^{\mathrm{F}} := \begin{bmatrix} \mathbf{u}^{\mathrm{F}}_{\mathrm{I}} \\ \mathbf{u}^{\mathrm{F}}_{\Gamma} \end{bmatrix} \quad \text{and} \quad \mathbf{f}^{\mathrm{F}} := \begin{bmatrix} \mathbf{f}^{\mathrm{F}}_{\mathrm{I}} \\ \mathbf{f}^{\mathrm{S}}_{\Gamma} + \mathbf{f}^{\mathrm{F}}_{\Gamma} \end{bmatrix},$$

the non-linear FSI problem (29) is rewritten as

$$\mathbf{f}^{\mathrm{FSI}}(\mathbf{x}^{\mathrm{FSI}}) = \mathbf{0} \quad \text{with } \mathbf{f}^{\mathrm{FSI}} := \begin{bmatrix} \mathbf{f}^{\mathrm{S}}_{\mathrm{I}} \\ \mathbf{f}^{\mathrm{G}}_{\mathrm{I}} \\ \mathbf{f}^{\mathrm{F}} \end{bmatrix} \text{ and } \mathbf{x}^{\mathrm{FSI}} := \begin{bmatrix} \mathbf{d}^{\mathrm{S}}_{\mathrm{I}} \\ \mathbf{d}^{\mathrm{G}}_{\mathrm{I}} \\ \mathbf{u}^{\mathrm{F}} \end{bmatrix}. \tag{30}$$

Following the approach by Küttler and Wall [55], the fully coupled FSI problem (30) is solved using a Newton-like method:

$$\mathbf{x}^{\mathrm{FSI}}_{(i+1)} = \mathbf{x}^{\mathrm{FSI}}_{(i)} + \Delta\mathbf{x}^{\mathrm{FSI}}_{(i)} \quad \text{with} \quad \left[\frac{\partial \mathbf{f}^{\mathrm{FSI}}}{\partial \mathbf{x}^{\mathrm{FSI}}}\right]_{(i)} \Delta\mathbf{x}^{\mathrm{FSI}}_{(i)} = -\mathbf{f}^{\mathrm{FSI}}_{(i)}. \tag{31}$$

Note that the Jacobian matrix $\partial\mathbf{f}^{\mathrm{FSI}}/\partial\mathbf{x}^{\mathrm{FSI}}$ involves the linearization of $\mathbf{f}^{\mathrm{S}}_{\mathrm{I}}$, $\mathbf{f}^{\mathrm{G}}_{\mathrm{I}}$, $\mathbf{f}^{\mathrm{F}}$ with respect to the variables $\mathbf{d}^{\mathrm{S}}_{\mathrm{I}}$, $\mathbf{d}^{\mathrm{G}}_{\mathrm{I}}$, $\mathbf{u}^{\mathrm{U}}$. This leads to a system of equations with the following block structure

$$\begin{bmatrix} \dfrac{\partial \mathbf{f}^{\mathrm{S}}_{\mathrm{I}}}{\partial \mathbf{d}^{\mathrm{S}}_{\mathrm{I}}} & \mathbf{0} & \dfrac{\partial \mathbf{f}^{\mathrm{S}}_{\mathrm{I}}}{\partial \mathbf{u}^{\mathrm{U}}} \\ \mathbf{0} & \dfrac{\partial \mathbf{f}^{\mathrm{G}}_{\mathrm{I}}}{\partial \mathbf{d}^{\mathrm{G}}_{\mathrm{I}}} & \dfrac{\partial \mathbf{f}^{\mathrm{G}}_{\mathrm{I}}}{\partial \mathbf{u}^{\mathrm{U}}} \\ \dfrac{\partial \mathbf{f}^{\mathrm{F}}}{\partial \mathbf{d}^{\mathrm{S}}_{\mathrm{I}}} & \dfrac{\partial \mathbf{f}^{\mathrm{F}}}{\partial \mathbf{d}^{\mathrm{G}}_{\mathrm{I}}} & \dfrac{\partial \mathbf{f}^{\mathrm{F}}}{\partial \mathbf{u}^{\mathrm{F}}} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{d}^{\mathrm{S}}_{\mathrm{I}} \\ \Delta\mathbf{d}^{\mathrm{G}}_{\mathrm{I}} \\ \Delta\mathbf{u}^{\mathrm{U}} \end{bmatrix} = - \begin{bmatrix} \mathbf{f}^{\mathrm{S}}_{\mathrm{I}} \\ \mathbf{f}^{\mathrm{G}}_{\mathrm{I}} \\ \mathbf{f}^{\mathrm{F}} \end{bmatrix}, \tag{32}$$

corresponding to the solid, grid motion and fluid fields.

Solving the system of equations (32) is very challenging because it combines physical fields of very different kind. Particular preconditioners for this problem are proposed for example in [57, 8, 58]. In the examples below, we compare the performance of the FSI-specific solvers by Gee et al. [8] with respect to the general purpose preconditioners and implementation proposed in Section 4. The considered methods are summarized in Table 2.

Name	Description of the preconditioner
BGS(AMG) generic	General purpose preconditioner given in Section 4.1.
AMG(BGS) generic	General purpose preconditioner given in Section 4.3 with BGS smoothers.
BGS(AMG) FSI	FSI-specific version of the preconditioner BGS(AMG) given in [8].
AMG(BGS) FSI	FSI-specific version of the preconditioner AMG(BGS) given in [8].

Table 2: Description of the preconditioners for fluid-structure interaction.

6.3 Lung model

The third and most challenging coupled problem considered in this paper is the lung model introduced by Yoshihara et al. [23, 24]. This model is an extension of a standard FSI problem in order to account for the inflation of the lung tissue in the respiration process. The key rationale of the model is the following.

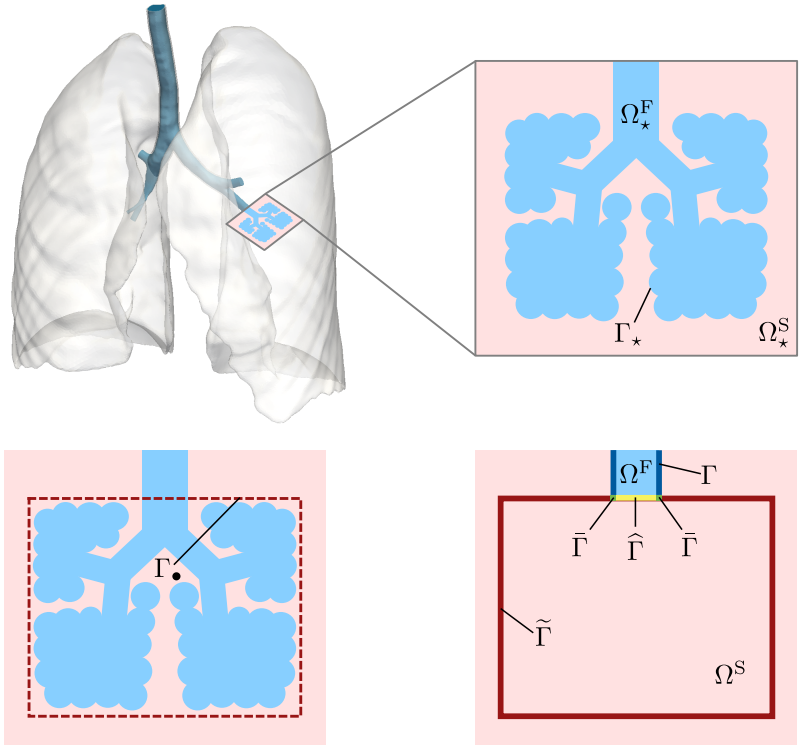

Figure 7: Illustration of the lung model introduced by Yoshihara et al. [23, 24].

The human lung can be ideally modeled as an FSI problem coupling the air transported by the respiratory tree and the lung tissue. Completely resolving this problem is usually not feasible due to the complex geometries involved: only few branches of the airway tree are imageable with CT-scans, and therefore, the domains $\Omega^{\mathrm{S}}_{\star}$, $\Omega^{\mathrm{F}}_{\star}$ and $\Gamma_{\star}$ of the ideal FSI problem are not available, see Figure 7 (top). In the computational model, the airways end in artificial outlets, but in reality, the branching of the conducting passages continues until the alveolar region is reached. The behavior of this region completely dominates the overall behavior of the system. Following the approach by Yoshihara et al. [23, 24], the unresolved structures are surrounded by an artificial boundary $\Gamma_{\bullet}$, see Figure 7 (bottom, left), and a suitable model is introduced inside the delimited region accounting for the unresolved features. The region enclosed by $\Gamma_{\bullet}$ is modeled as an elastic solid with the same properties as the surrounding lung tissue, see Figure 7 (bottom, right). The discarded fluid domain is taken into account by assuming that the air transported in the feeding vessel ending at $\widehat{\Gamma}$ causes an inflation of the region delimited by $\Gamma_{\bullet}$. That is, the volume change of the region enclosed by $\Gamma_{\bullet}$ is equal to the volume of air flowing into it through $\widehat{\Gamma}$. This constraint

is expressed formally as

$$\dot{V}_{\Gamma_\bullet} = \int_{\widehat{\Gamma}} (\boldsymbol{u}^{\mathrm{F}} - \dot{\boldsymbol{d}}^{\mathrm{G}}) \cdot \boldsymbol{n} \, \mathrm{d}A, \tag{33}$$

where $\dot{V}_{\Gamma_\bullet}$ denotes the rate of change of the volume surrounded by $\Gamma_\bullet$ and the right hand side is the flow rate through the outlet $\widehat{\Gamma}$.

In order to have a well defined problem, extra conditions have to be introduced on the artificial ends of the airways:

$$\boldsymbol{u}^{\mathrm{F}} = \dot{\boldsymbol{d}}^{\mathrm{S}} \quad \text{on } \bar{\Gamma}, \tag{34a}$$
$$\boldsymbol{d}^{\mathrm{G}} = \boldsymbol{d}^{\mathrm{S}} \quad \text{on } \bar{\Gamma} \cup \widehat{\Gamma}, \tag{34b}$$
$$\boldsymbol{\sigma}^{\mathrm{F}} \cdot \boldsymbol{n} = \boldsymbol{0} \quad \text{on } \widehat{\Gamma}, \tag{34c}$$
$$\boldsymbol{\sigma}^{\mathrm{S}} \cdot \boldsymbol{n} = \boldsymbol{0} \quad \text{on } \widehat{\Gamma}. \tag{34d}$$

The first condition is the standard continuity of velocities inherited from the FSI interface. The second condition is introduced in order to avoid gaps between the fluid and solid domains on the artificial boundary $\bar{\Gamma} \cup \widehat{\Gamma}$, and the two last are standard "do nothing" conditions.

In conclusion, the lung model results in an FSI problem which is augmented with the constraint (33) and the extra conditions (34). This leads (after the usual space and time discretization) to a non-linear system of algebraic equations to be solved at each time step. Defining this problem requires introducing the discrete version of constraint (33), which is denoted here using the scalar equation

$$c(\mathbf{d}^{\mathrm{S}}, \mathbf{u}^{\mathrm{F}}) = 0.$$

The number of such constraints in the model is equal to the number of artificial outlets $\widehat{\Gamma}$. For convenience, the constraints corresponding to all the artificial outlets are grouped in the vector equation

$$\mathbf{c}(\mathbf{d}^{\mathrm{S}}, \mathbf{u}^{\mathrm{F}}) = 0. \tag{35}$$

With this notation, the non-linear algebraic equations to be solved at each time step are

$$\mathbf{f}^{\mathrm{C;FSI}}(\mathbf{x}^{\mathrm{C;FSI}}) = \mathbf{0},$$

where

$$\mathbf{f}^{\mathrm{C;FSI}} = \begin{bmatrix} \mathbf{f}^{\mathrm{S}}_{\mathrm{I}\cup\widehat{\Gamma}} \\ \mathbf{f}^{\mathrm{G}}_{\mathrm{I}} \\ \mathbf{f}^{\mathrm{F}}_{\mathrm{I}\cup\widehat{\Gamma}} \\ \mathbf{f}^{\mathrm{S}}_{\Gamma\cup\bar{\Gamma}} + \mathbf{f}^{\mathrm{F}}_{\Gamma\cup\bar{\Gamma}} \\ \mathbf{c} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\lambda}^{\mathrm{T}} \dfrac{\partial \mathbf{c}}{\partial \mathbf{d}^{\mathrm{S}}_{\mathrm{I}\cup\widehat{\Gamma}}} \\ \boldsymbol{\lambda}^{\mathrm{T}} \dfrac{\partial \mathbf{c}}{\partial \mathbf{d}^{\mathrm{G}}_{\mathrm{I}}} \\ \boldsymbol{\lambda}^{\mathrm{T}} \dfrac{\partial \mathbf{c}}{\partial \mathbf{u}^{\mathrm{F}}_{\mathrm{I}\cup\widehat{\Gamma}}} \\ \boldsymbol{\lambda}^{\mathrm{T}} \dfrac{\partial \mathbf{c}}{\partial \mathbf{u}^{\mathrm{F}}_{\Gamma\cup\bar{\Gamma}}} \\ \mathbf{0} \end{bmatrix} \quad \text{and} \quad \mathbf{x}^{\mathrm{C;FSI}} = \begin{bmatrix} \mathbf{d}^{\mathrm{S}}_{\mathrm{I}\cup\widehat{\Gamma}} \\ \mathbf{d}^{\mathrm{G}}_{\mathrm{I}} \\ \mathbf{u}^{\mathrm{F}}_{\mathrm{I}\cup\widehat{\Gamma}} \\ \mathbf{u}^{\mathrm{F}}_{\Gamma\cup\bar{\Gamma}} \\ \boldsymbol{\lambda} \end{bmatrix}. \tag{36}$$

Note that the first four equations are essentially the standard FSI equations (29) extended with the interface conditions (34) and the loads introduced by the Lagrange multipliers $\boldsymbol{\lambda}$ associated with the constraint (35). The last equation in (36) corresponds to the algebraic constraint $\mathbf{c}$. After linearization of problem (36) within a Newton-like procedure, the system to be solved at each Newton step has the

form:

$$\left[\begin{array}{ccc|c} \mathbf{A}^{\mathrm{SS}} & \mathbf{A}^{\mathrm{SG}} & \mathbf{A}^{\mathrm{SF}} & \mathbf{A}^{\mathrm{SC}} \\ \mathbf{A}^{\mathrm{GS}} & \mathbf{A}^{\mathrm{GG}} & \mathbf{A}^{\mathrm{GF}} & \mathbf{A}^{\mathrm{GC}} \\ \mathbf{A}^{\mathrm{FS}} & \mathbf{A}^{\mathrm{FG}} & \mathbf{A}^{\mathrm{FF}} & \mathbf{A}^{\mathrm{FC}} \\ \hline \mathbf{A}^{\mathrm{CS}} & \mathbf{A}^{\mathrm{CG}} & \mathbf{A}^{\mathrm{CF}} & \mathbf{0} \end{array}\right] \left[\begin{array}{c} \Delta \mathbf{d}^{\mathrm{S}}_{\mathrm{I}\cup\widehat{\Gamma}} \\ \Delta \mathbf{d}^{\mathrm{G}}_{\mathrm{I}} \\ \Delta \mathbf{u}^{\mathrm{F}} \\ \hline \Delta \boldsymbol{\lambda} \end{array}\right] = - \left[\begin{array}{c} \mathbf{b}^{\mathrm{S}} \\ \mathbf{b}^{\mathrm{G}} \\ \mathbf{b}^{\mathrm{F}} \\ \hline \boldsymbol{c} \end{array}\right]. \tag{37}$$

This system of equations consists of 4 matrix equations associated with the 3 FSI fields (solid, grid motion and fluid) plus a block associated with the constraint. The FSI blocks are essentially the same as the ones already presented in the FSI system (32). The precise definition of all the objects in (37) is not provided here for brevity, refer to Yoshihara et al. [23, 24] for details.

Solving the system of equations (37) is very challenging. In the one hand, it shares the difficulties already commented for the pure FSI system (32). Moreover, the system has a zero matrix on the main diagonal due to the saddle point structure introduced by the volumetric constraint (35). This means that the BGS method cannot be used alone for solving the system, and therefore, the preconditioner BGS(AMG) introduced in Section 4.1 cannot be applied.

The preconditioners considered here for the lung model are based on the SIMPLE method (cf. Section 4.2). Recall that in order to use SIMPLE, the 4×4 block structure of system (37) has to be transformed into a 2×2 structure. This is achieved here by grouping the FSI blocks as indicated in equation (37). The SIMPLE requires solving systems with the predictor matrix $\mathbf{A}_{11}$ and Schur complement matrix $\widetilde{\mathbf{S}}$. For the splitting indicated in equation (37), the matrix $\mathbf{A}_{11}$ is

$$\mathbf{A}_{11} = \left[\begin{array}{ccc} \mathbf{A}^{\mathrm{SS}} & \mathbf{A}^{\mathrm{SG}} & \mathbf{A}^{\mathrm{SF}} \\ \mathbf{A}^{\mathrm{GS}} & \mathbf{A}^{\mathrm{GG}} & \mathbf{A}^{\mathrm{GF}} \\ \mathbf{A}^{\mathrm{FS}} & \mathbf{A}^{\mathrm{FG}} & \mathbf{A}^{\mathrm{FF}} \end{array}\right], \tag{38}$$

corresponding to the FSI operator. That is, at each iteration of the SIMPLE method, a linearized FSI problem has to be solved. To solve this problem we consider the two FSI preconditioners previously presented in Table 2, namely BGS(AMG), and AMG(BGS).

On the other hand, the system involving the Schur complement $\widetilde{\mathbf{S}}$ is solved with a direct solver because the size of $\widetilde{\mathbf{S}}$ is frequently rather small in this application. This is because, the number of rows in $\widetilde{\mathbf{S}}$ corresponds to the number of ending outlets in the lung model which in practice is much smaller than the number of degrees of freedom in the other fields.

In conclusion, we consider two different preconditioners for the lung model, see Table 3. The first one is denoted as SIMPLE(BGS(AMG)) and consists in an outer SIMPLE iteration (Section 4.2) and an inner BGS(AMG) preconditioner (cf. Section 4.1) for solving the FSI block (38). The second preconditioner is denoted as SIMPLE(AMG(BGS)) and is an outer SIMPLE iteration and an inner monolithic AMG preconditioner (cf. Section 4.3) for solving the FSI block. The only difference between these to methods it that SIMPLE(BGS(AMG)) resolves the FSI coupling only at the finest AMG level, while SIMPLE(AMG(BGS)) resolves the the FSI coupling at all levels, see Figure 8.

Name	Description of the preconditioner
SIMPLE(BGS(AMG))	Outer SIMPLE iteration using an inner BGS(AMG) preconditioner for the FSI subproblem.
SIMPLE(AMG(BGS))	Outer SIMPLE iteration using an inner AMG(BGS) preconditioner for the FSI subproblem.

Table 3: Description of the preconditioners for the lung model.

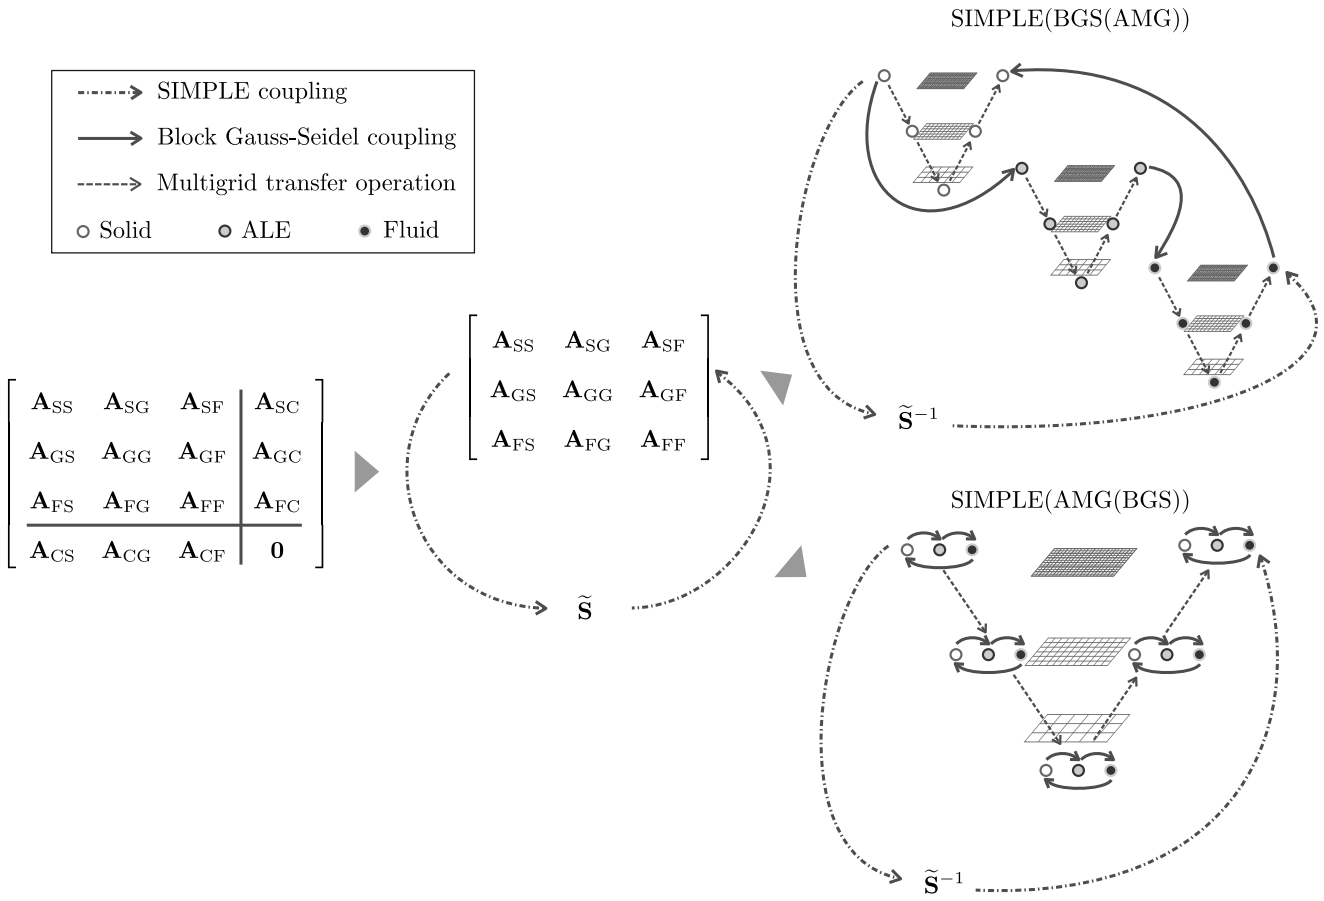

Figure 8: Illustration of the preconditioners for the lung model. We consider two different methods, namely SIMPLE(BGS(AMG)) and SIMPLE(AMG(BGS). Both preconditioners are obtained by merging the fields corresponding to the FSI problem to recover a 2×2 block structure in order to apply SIMPLE. Within the SIMPLE iteration, a linearized FSI problem has to be solved. We consider two options, either a BGS(AMG) method enforcing the coupling at the finest multigrid level, or a AMG(BGS) method resolving the coupling at all levels.

7 Numerical examples

7.1 TSI example: Second Danilovskaya problem

This first example is devoted to study the performance and scalability of the proposed preconditioners in a benchmark test for TSI. The example is known as the *second Danilovskaya problem*, it was originally proposed in [59], and it has often been used in literature for validating fully coupled thermo-mechanical models (e.g. in [7, 53, 60]).

The example consists in a thermo-elastic rectangular prism which is heated on the top face, see Figure 9. The thermo-mechanical properties of the material are Young's modulus $E = 210 \cdot 10^9$ Pa, Poisson's ratio $\nu = 0.3$, density $\rho = 7.86 \cdot 10^3$ kg/m^3, heat capacity $C = 0.821$ J/(kg $\cdot$ K), heat conductivity $k = 1.03$ W/(m $\cdot$ K) and coefficient of thermal expansion $\alpha = 1.1 \cdot 10^{-5}$ 1/K. The body is initially at rest, it has an initial temperature $u_0 = 273.15$ K, it is clamped at the bottom face, and it is heated at the top face with an ambient temperature $u_\infty = 372.15$ K. The external heating introduced by the ambient temperature u_∞ is modeled as a prescribed flux,

$$\boldsymbol{q}(u) \cdot \boldsymbol{n} = h(u - u_\infty), \tag{39}$$

proportional to the difference between the current surface temperature u and the ambient temperature u_∞. The boundary condition (39) is the so-called heat convection boundary condition, and the proportionally constant h is the so-called linear heat transfer coefficient. Its value is $h = \rho C \bar{h}$, with $\bar{h} = 1 \cdot 10^{-5}$

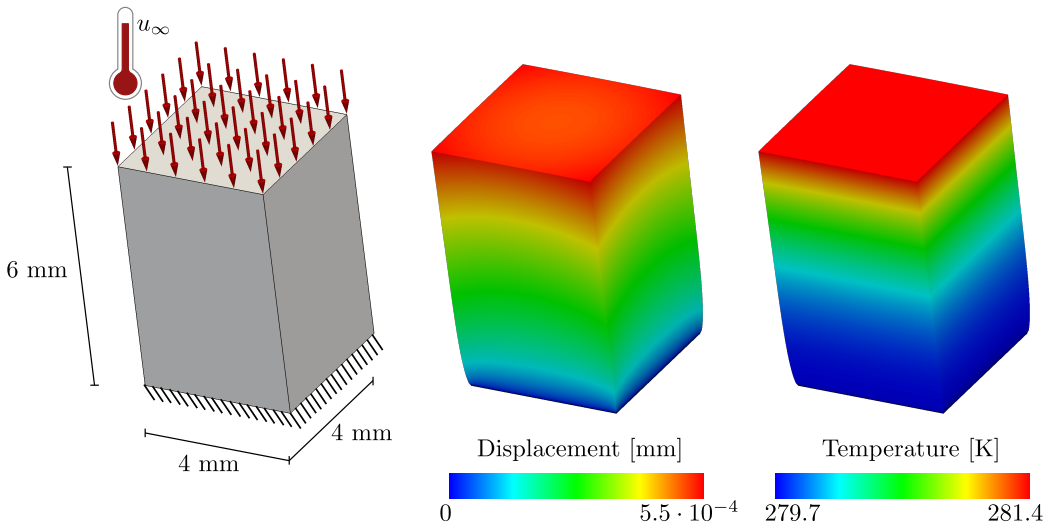

Figure 9: Second Danilovskaya problem: Problem description (left), deformation (center) and temperature distribution (right) at time $t = 4$ s.

m/s. The other faces of the body are thermally insulated, i.e. homogeneous thermal Neumann conditions are applied. The problem is discretized in space with trilinear hexahedral finite elements. Four different meshes are considered in this example, see Table 4. The time discretization is carried out with the one step-θ method with $\theta = 2/3$ and time step $\Delta t = 4 \cdot 10^{-2}$ s.

Mesh Id	Processors	Structural field	Thermal field	Total	Total/Processors
1	2	63888	21296	85184	42592
2	8	235824	78608	314432	39304
3	32	998250	332750	1331000	41593
4	128	3951018	1317006	5268024	41156

Table 4: Second Danilovskaya problem: Number of processors and number of degrees of freedom for the four meshes considered in the example.

The discrete non-linear problems to be solved at each time step are solved with the monolithic Newton iteration (24) and the resulting linear systems are solved using GMRES preconditioned with the methods AMG(BGS), BGS(AMG) and SIMPLE(AMG) given in Table 1. For the outer Newton iterations, convergence is declared when the full residual fulfills $\|\mathbf{f}^{\mathrm{TSI}}\|_{\mathrm{rms}} < 10^{-6}$ and the residual associated with the individual structural and thermal fields fulfill $\|\mathbf{f}^{\mathrm{S}}\|_{\mathrm{rms}} < 10^{-8}$ and $\|\mathbf{f}^{\mathrm{Th}}\|_{\mathrm{rms}} < 10^{-8}$, where $\|\cdot\|_{\mathrm{rms}}$ is the root mean square norm. The inner GMRES iteration is stopped when the residual of the linear problem is such that $\|\mathbf{r}\|_2/\|\mathbf{r}^0\|_2 < 10^{-8}$, where $\mathbf{r}^0$ stands for the residual associated with the initial guess given to the GMRES solver. Using these tolerances, three Newton iterations are required for achieving convergence at each time step.

The following parameters are considered for the studied preconditioners. The number of iterations for BGS and SIMPLE is taken as $k_{\max} = 1$. The underlying single field AMG methods are built using standard smoothed aggregation AMG [40, 17] available in the MueLu [41] library. The multigrid smoothers $\boldsymbol{S}_i^\ell$ for the single field problems are defined as an additive Schwartz domain decomposition using a damped Gauss-Seidel iteration (damping $\omega = 0.79$) as sub-domain solver. The coarse level solver is the so-called KLU direct solver given in the Amesos package [61].

Figure 10 shows the average number of linear solver iterations per Newton step and the associated total CPU time (including the setup of the preconditioners) for the three preconditioners considered here and for the finest mesh in Table 4. The results show that the methods lead to a reasonable amount of GMRES

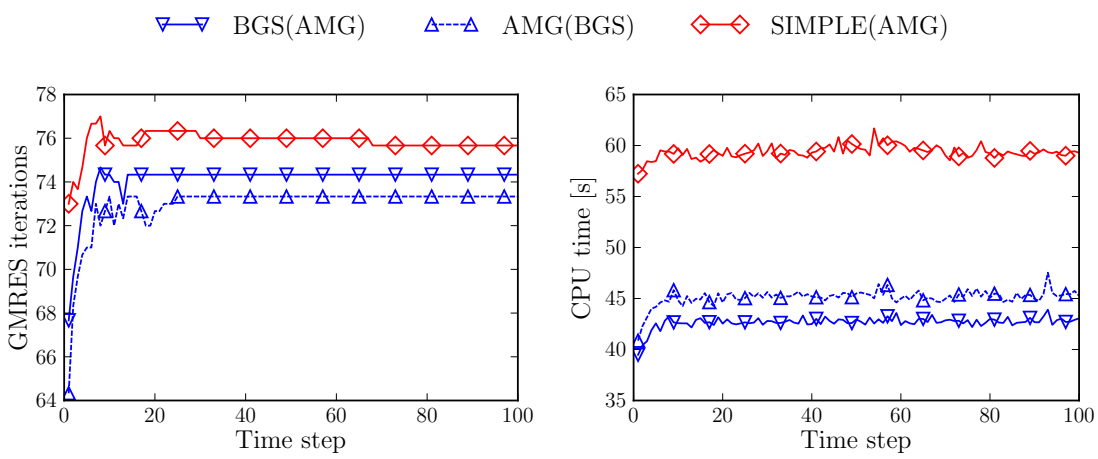

Figure 10: Second Danilovskaya problem: Average of linear solver iterations per Newton step (left) and the corresponding CPU times (right). The results are for the finest discretization given in Table 4 (5.2 milion unknowns on 128 processors).

iterations for this very fine mesh which means that the proposed methods are preconditioning the problem effectively. The best preconditioner in terms of iteration count is the AMG(BGS) method. This is what we expected since AMG(BGS) resolves the coupling at all multigrid levels. However, the BGS(AMG) method is clearly better in terms of CPU times, which shows that the benefit of using AMG(BGS) (which is a more complex method) is not clearly justified for this example and the chosen parameters. On the other hand, the SIMPLE(AMG) method is the worst in terms of iterations and CPU time. It might seem that SIMPLE(AMG) is not a suitable preconditioner for this type of problems. However, we show later in this example that for strong values of the thermal expansion coefficient α, SIMPLE(AMG) is performing better than the other two methods (see Figure 12).

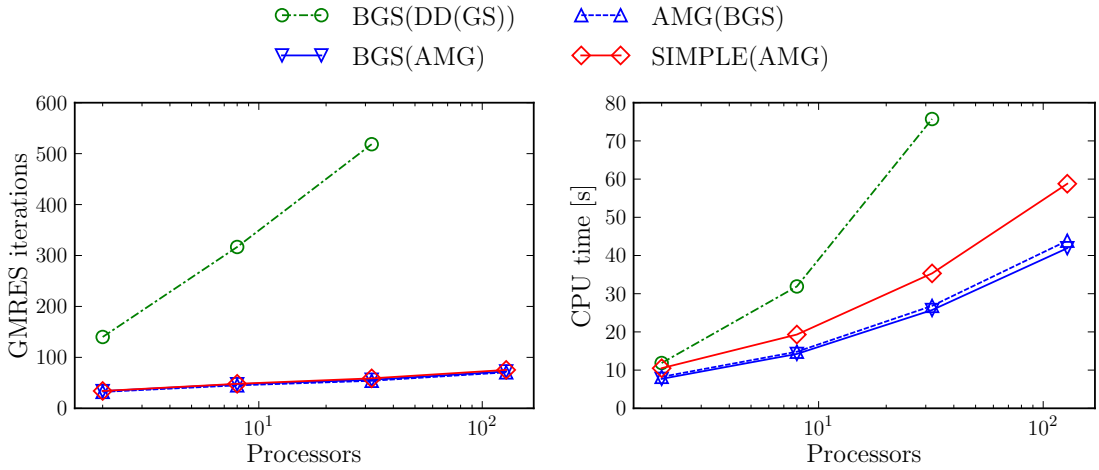

Figure 11: Second Danilovskaya problem: Results of the weak scalability study. The CPU times include the setup costs of the preconditioner.

Figure 11 and Tables 5 and 6 report a weak scalability test. The test is done using four meshes given in Table 4 and setting the number of processors such that the work per process is fairly the same in all cases. The proposed multigrid methods AMG(BGS), BGS(AMG) and SIMPLE(AMG) are compared with a standard single-level preconditioner denoted here as BGS(DD(GS)). This single-level method is defined as an outer BGS iteration with standard additive Schwartz domain decomposition (DD) methods for solving the resulting uncoupled problems. The local subdomain solvers in the domain decomposition are damped Gauss-Seidel (GS) iterations with damping $\omega = 0.79$. The single level method is one of

the most simple parallel preconditioners that can be considered in such a problem, and therefore, it is considered here as a reference.

Processors	D.O.F.	BGS(DD(GS))	BGS(AMG)	AMG(BGS)	SIMPLE(AMG)
2	85184	140	33	32	34
8	314432	316	47	45	48
32	1331000	518	56	54	58
128	5268024	–	72	70	75

Table 5: Second Danilovskaya problem: Average of linear solver iterations in the weak scalability test reported in Figure 11.

Processors	D.O.F.	BGS(DD(GS))	BGS(AMG)	AMG(BGS)	SIMPLE(AMG)
2	85184	11.8	7.7	8.2	10.4
8	314432	31.9	14.2	14.8	19.3
32	1331000	75.7	25.7	26.8	35.3
128	5268024	–	41.9	43.8	58.7

Table 6: Second Danilovskaya problem: CPU times in the weak scalability test reported in Figure 11.

Figure 11 and Tables 5 and 6 show that the presented multigrid methods AMG(BGS), BGS(AMG) and SIMPLE(AMG) lead to iteration counts and CPU times that mildly depend on the problem size. Obviously, this is not the optimal multigrid behavior achieved in single-field elliptic problems (i.e. independence of the iteration count and CPU times with respect to the problem size), but it can be considered a very good scalability if one takes into account that the optimal multigrid behavior is rarely reproduced in complex multiphysics. On the other hand, it is observed that the single-level method BGS(DD(GS)) is not scalable since the iteration count strongly depends on the problem size. This is what we expected since it is well known that single-level domain decomposition methods does not scale properly (cf. [13]). In conclusion, this scalability study shows that proposed preconditioners AMG(BGS), BGS(AMG) and SIMPLE(AMG) have a decent scalability, and therefore, can be considered to solve this thermo-mechanical problem efficiently in parallel.

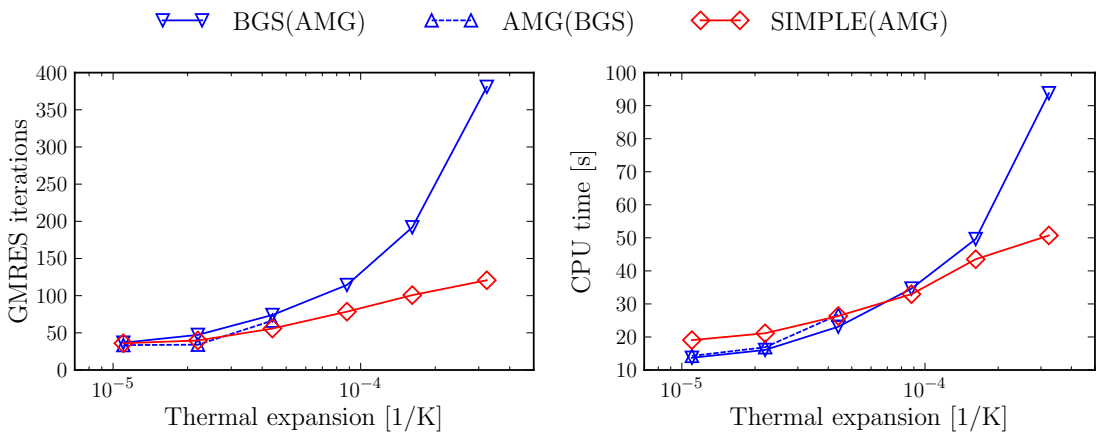

Figure 12: Second Danilovskaya problem: Dependence of the performance of the preconditioners with respect to the amount of thermo-mechanical coupling. The results correspond to the first mesh given in Table 4.

Finally, we investigate the performance of the preconditioners with respect to the thermal expansion coefficient α, which controls the coupling between the structural and thermal fields (the coupling becomes

stronger when α increases). As it is reported in Figure 12, the performance of the methods AMG(BGS) and BGS(AMG) strongly degrades when the thermal expansion coefficient α increases. This is explained by noting that the off-diagonal blocks of the system matrix in equation (25) are proportional to α. Thus, when α grows, the off-diagonal blocks discarded by the BGS iteration become dominant and affect negatively the performance of the method. The AMG(BGS) preconditioner is more sensitive to this effect than BGS(AMG). Note that for AMG(BGS), the convergence of the linear solver is achieved only up to the third value of the thermal expansion coefficient $\alpha = 4.4 \cdot 10^{-5}$ 1/K. On the other hand, the SIMPLE(AMG) method is the most robust technique with respect the thermal expansion coefficient α. In conclusion, for mild values of the thermal expansion coefficient, the BGS(AMG) method is preferred, whereas for strong values, the best choice is SIMPLE(AMG).

7.2 TSI example: Rocket nozzle

The aim of this second example is to study the performance of the proposed preconditioners in a real-world problem. The example is taken form [7] and simulates the thermo-mechanical behavior of a rocked nozzle. The nozzle geometry (see Figure 13) is inspired by the Vulcain rocket engine installed in the Ariane space launcher. The nozzle is 0.3985 m in height and its cross section is approximated by a circular ring. The inner radius at the nozzle's inlet is 0.0253 m and 0.1698 m at the outlet. The nozzle is equipped with 81 equidistant cooling channels which traverse the body longitudinally. Thanks to the symmetry of the problem, only a 1/81-th of the nozzle is meshed and simulated by introducing the required symmetry conditions. The computational domain is a sector of the nozzle of angle 360°/81 which includes only one cooling channel, see Figure 13 (center).

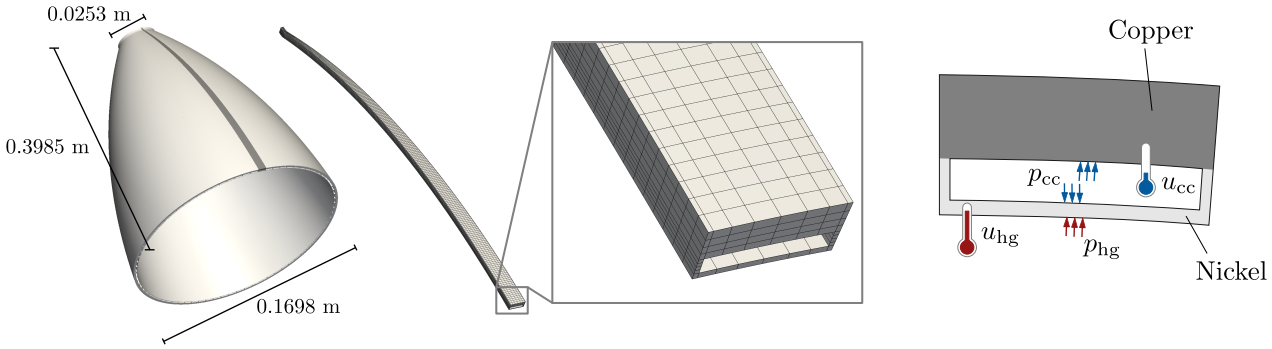

Figure 13: Rocket nozzle example: Full geometry of the nozzle (left), computational domain including one cooling channel (center), and generic cross section of the computational domain with the applied thermo-mechanical loads (right).

The rocket nozzle is made of two materials: an external copper alloy body and an internal nickel jacket, see Figure 13 (right). The copper alloy has Young's modulus $E = 148$ GPa, Poisson's ratio $\nu = 0.3$, density $\rho = 9.13 \cdot 10^3$ kg/m^3, heat conductivity $k = 310$ W/(m $\cdot$ K), heat capacity $C = 373$ J/(kg $\cdot$ K) and coefficient of thermal expansion $\alpha = 1.72 \cdot 10^{-5}$ 1/K. On the other hand, the nickel jacked has the following properties: $E = 193$ GPa, $\nu = 0.3$, $\rho = 8910$ kg/m^3, $k = 75$ W/(m $\cdot$ K), $C = 444$ J/(kg $\cdot$ K) and $\alpha = 1.22 \cdot 10^{-5}$ 1/K.

The rocked nozzle is initially at rest and cooled down with a uniform initial temperature of 40 K. The thermo-mechanical loads acting on the body are the high pressure and temperature of the exhausted hot gases (namely p_{hg} and u_{hg}) and the pressure and low temperature of the fluid filling the cooling channels (namely p_{cc} and u_{cc}). The temporal evolution of p_{hg}, u_{hg}, p_{cc} and u_{cc} is taken from [62] corresponding to a prototypical loading cycle for rocked nozzles, see Figure 14. As in the previous example, the effect of the external temperatures u_{hg} and u_{cc} is modeled with the heat convection boundary condition (39). The linear heat transfer coefficient h is $h_{hg} = 32$ kW/(m^{2}K) for the hot gas, and $h_{cc} = 100$ kW/(m^{2}K) for the cooling channel. Additionally, the inlet of the nozzle is fixed in longitudinal direction in order to

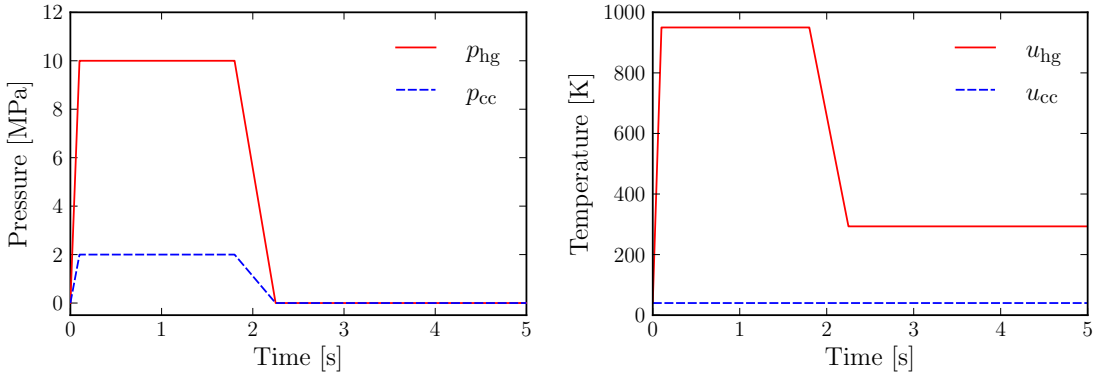

Figure 14: Rocket nozzle example: Temporal evolution of the thermo-mechanical loads (pressure and temperature) associated with the exhausted hot gas (p_{hg}, u_{hg}) and with the liquid filling the cooling channels (p_{cc}, u_{cc}).

take into account the combustion chamber which is attached to it. The total simulation time is $t = 5$ s. A screen-shot of the deformation of the nozzle and the temperature field is given in Figure 15.

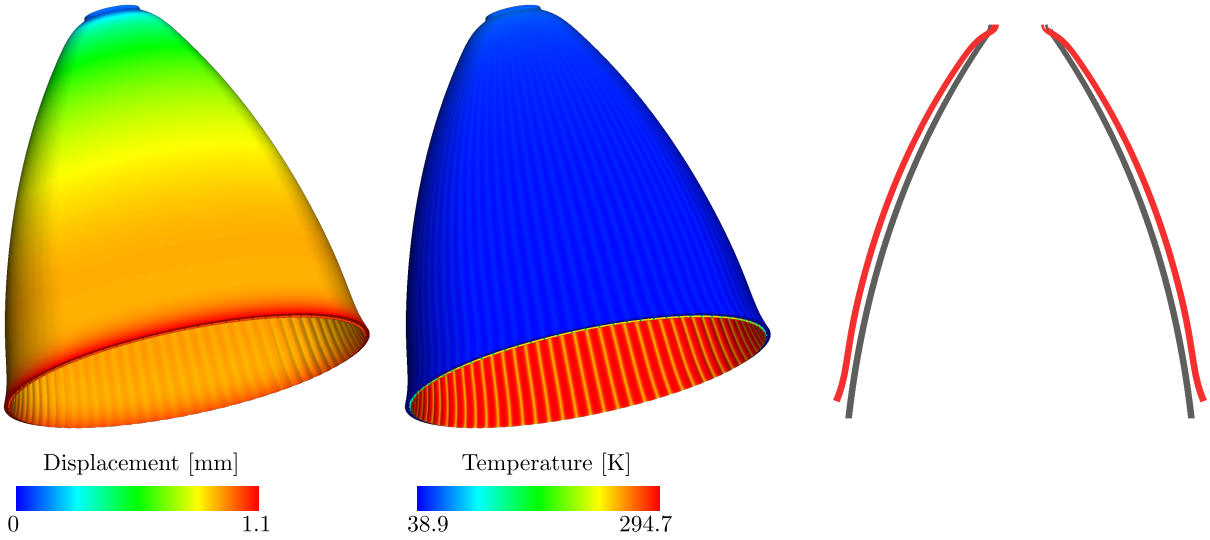

Figure 15: Rocket nozzle example: Deformation of the nozzle (left), temperature distribution (center) and original and deformed longitudinal section (right). The results are given at time $t = 1$ s and the deformation is magnified 20 times.

The performance of the preconditioners is studied with a weak scalability test. To this end, several finite element meshes are considered and the number of processors is chosen such that the work per processor is similar for all the meshes, see Table 7. As in the previous example, tri-linear hexahedral finite elements are used for the space discretization and the one step-θ with $\theta = 2/3$ for the time discretization. The time step is chosen as $\Delta t = 5 \cdot 10^{-2}$ s. The Newton iterations are stopped when the full non-linear residuum is $\|\mathbf{f}^{\text{TSI}}\|_{\text{rms}} < 10^{-6}$, and the entries of the non-linear residuum corresponding to the structural and thermal fields are $\|\mathbf{f}_{\text{S}}\|_{\text{rms}} < 10^{-6}$ and $\|\mathbf{f}_{\text{T}}\|_{\text{rms}} < 10^{-6}$ respectively. The GMRES iteration is stopped when the linear residuum fulfills $\|\mathbf{r}\|_2/\|\mathbf{r}^0\|_2 < 10^{-7}$. Using these tolerances, between 2 and 3 Newton iterations are required for achieving convergence at each time step. The preconditioners AMG(BGS), BGS(AMG), SIMPLE(AMG) and BGS(DD(GS)) are the same as in the previous example.

The results of the weak scalability study are reported in Figure 16 and in Tables 8 and 9. The performance of the preconditioners in this complex three dimensional example is similar to the behavior observed in the previous benchmark test (Section 7.1). Again, the multigrid methods AMG(BGS), BGS(AMG) and SIMPLE(AMG) have a very good scalability as the iteration count and CPU time of the linear solver

Mesh Id	Processors	Solid	Thermal field	Total	Total/Processors
1	4	47025	15675	62700	15675
2	32	307046	102582	410328	12822
3	256	2189124	729708	2918832	11401

Table 7: Rocket nozzle example: Number of processors and number of degrees of freedom for the three meshes considered in the example.

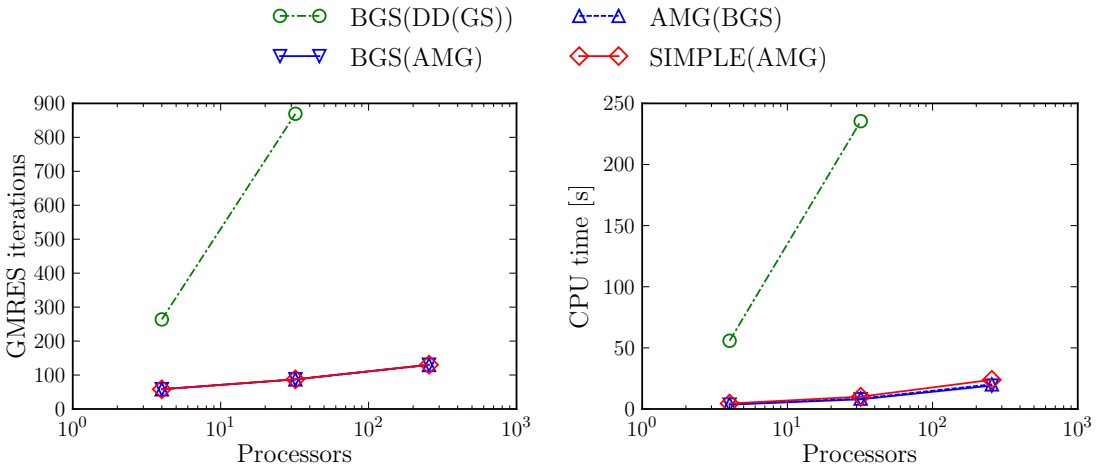

Figure 16: Rocket nozzle example: Results of the weak scalability study. The CPU times include the setup costs of the preconditioner.

Processors	D.O.F.	BGS(DD(GS))	BGS(AMG)	AMG(BGS)	SIMPLE(AMG)
4	62700	263	58	57	58
32	410328	869	87	87	87
256	2918832	–	130	129	130

Table 8: Rocket nozzle example: Average of linear solver iterations in the weak scalability test reported in Figure 16.

Processors	D.O.F.	BGS(DD(GS))	BGS(AMG)	AMG(BGS)	SIMPLE(AMG)
4	62700	55.7	3.5	3.9	4.5
32	410328	235.4	7.9	8.5	10.1
256	2918832	–	19.1	20.1	23.9

Table 9: Rocket nozzle example: CPU times in the weak scalability test reported in Figure 16.

increase only mildly with the problem size. On the other hand, the single-level method BGS(DD(GS)) based on a standard additive Schwarz domain decomposition is not a scalable method as previously observed. Note that the performance of the single-level method BGS(DD(GS)) is particularly poor in this complex example which demonstrates that the multigrid preconditioners AMG(BGS), BGS(AMG) and SIMPLE(AMG) are doing a particularly efficient job in this challenging setting. In conclusion, the multigrid methods AMG(BGS), BGS(AMG) and SIMPLE(AMG) have a good scalability also for this challenging example, which shows that the methods are able to solve this complex real-world problem efficiently.

7.3 FSI example: Pressure wave

The aim of this example is to compare the performance of the proposed general purpose preconditioners with respect to the FSI-specific preconditioners proposed by Gee et al. [8]. The study is carried out for the well known benchmark test introduced by Gerbeau and Vidrascu [63] consisting in a fluid wave traveling through a flexible tube which mimics hemodynamic conditions, see Figure 17. The tube has a length of $5 \cdot 10^{-2}$ m, an inner radius of $5 \cdot 10^{-3}$ m and an outer radius of $6 \cdot 10^{-3}$ m. The structural model used for the tube is a geometrically non-linear Saint Venant-Kirchhoff material with Young's modulus $E = 3 \cdot 10^5$ Pa, Poisson's ratio $\nu = 0.3$ and density $\rho^{\mathrm{s}} = 1.2 \cdot 10^3$ Kg/m^3. The tube is filled with a Newtonian fluid with density $\rho^{\mathrm{f}} = 1 \cdot 10^3$ Kg/m^3 and viscosity $\mu = 3 \cdot 10^{-3}$ Pa $\cdot$ s. Both the tube and the contained fluid are at rest at the initial time. The tube is clamped at both ends and the fluid is initially loaded with a surface traction of $1.33 \cdot 10^3$ Pa on one side for a time interval of $3 \cdot 10^{-3}$ s. This action results in a pressure wave traveling through the tube, see Figure 17.

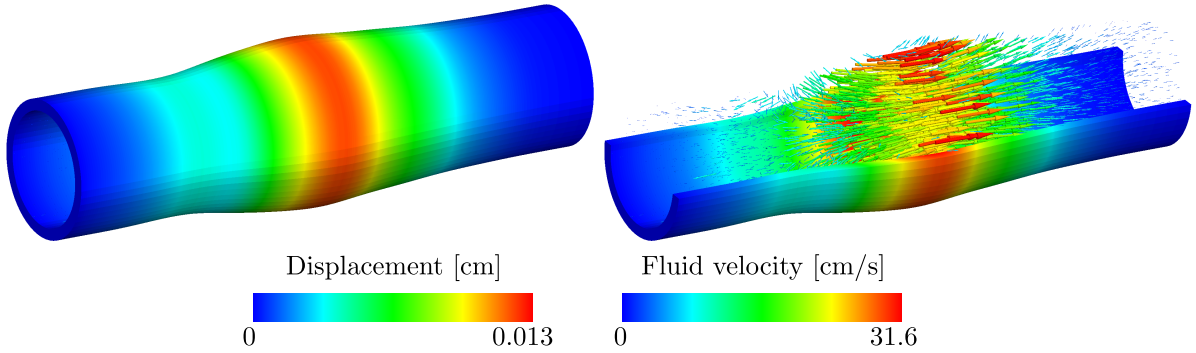

Figure 17: FSI example: Pressure wave traveling through a flexible tube. The deformation is magnifyied 10 times.

The problem is discretized in space using four different finite element meshes, see Table 10. Trilinear hexahedral elements are used in the structural domain and the fluid is discretized with stabilized trilinear hexahedral elements using the same interpolation for both velocity and pressure. The time discretization is performed using the trapezoidal rule for the structural domain and with the one step-θ method for the fluid ($\theta = 2/3$). The time step is $\Delta t = 10^{-4}$ s.

Mesh Id	Processors	Solid	Fluid	Grid	Total	Total/Processors
1	1	7560	23660	13965	45185	45185
2	4	24708	97128	60492	182328	45582
3	16	95580	372880	247800	716260	44766
4	64	396000	1434752	996864	2827616	44181

Table 10: FSI example: Number of processors and number of degrees of freedom for the four meshes considered in the example.

The non-linear FSI problems associated with each time step are solved with the monolithic Newton iteration (31) and using a preconditioned GMRES method for solving the resulting linear systems. The outer Newton iteration is stopped when the full FSI-residual fulfills $\|\mathbf{f}^{\mathrm{FSI}}\|_{\mathrm{rms}} < 10^{-6}$ and $\|\mathbf{f}^{\mathrm{FSI}}\|_{\infty} < 10^{-6}$ and the linear solver is stopped when the linear residual fulfills $\|\mathbf{r}\|_2 / \|\mathbf{r}^0\|_2 < 10^{-10}$.

The preconditioners studied here (Table 2) require building AMG methods for each of the FSI fields. For the structural field and the ALE field, the underlying multigrid hierarchies are built using standard smoothed aggregation AMG [40, 17]. For the fluid field, a Petrov-Galerkin [64] AMG method is considered which is specially designed for the non-symmetric structure of the fluid equations. The multigrid smoothers are chosen as an additive Schwartz domain decomposition using a damped Gauss-Seidel iteration as subdomain solver, see Table 11. This smoother is referred to as DD(GS(ω)), where ω is the

damping factor of the Gauss-Seidel iteration. The coarse level solver is the KLU direct solver given in the Amesos package [61].

Field	Smoothers	Coarsest level solver	Transfer operators	# of levels
Solid	DD(GS(0.79))	KLU	Smoothed aggregation AMG	3
Fluid	DD(GS(0.69))	KLU	Petrov-Galerkin AMG	3
ALE	DD(GS(0.79))	KLU	Smoothed aggregation AMG	3

Table 11: FSI example: Details of the AMG methods used to build the preconditioners.

Figure 18 shows the performance of the general purpose preconditioners BGS(AMG) and AMG(BGS) and their corresponding FSI-specific implementations proposed by Gee et al. [8]. Note that both implementations lead to similar CPU time which show that using a generic implementation (which is not optimized for this particular problem) is not affecting the performance of the solver. In fact, the generic implementation has slightly better CPU times, which can be due to the fact that different packages are considered for building the underlying AMG methods: the FSI-specific implementation uses ML [42] and the general implementation proposed here uses MueLu [41] which is a more recent multigrid package. This, shows that the performance of the methods is driven by the chosen AMG libraries (which are black boxes here), and not by the implementation of the coupling iteration.

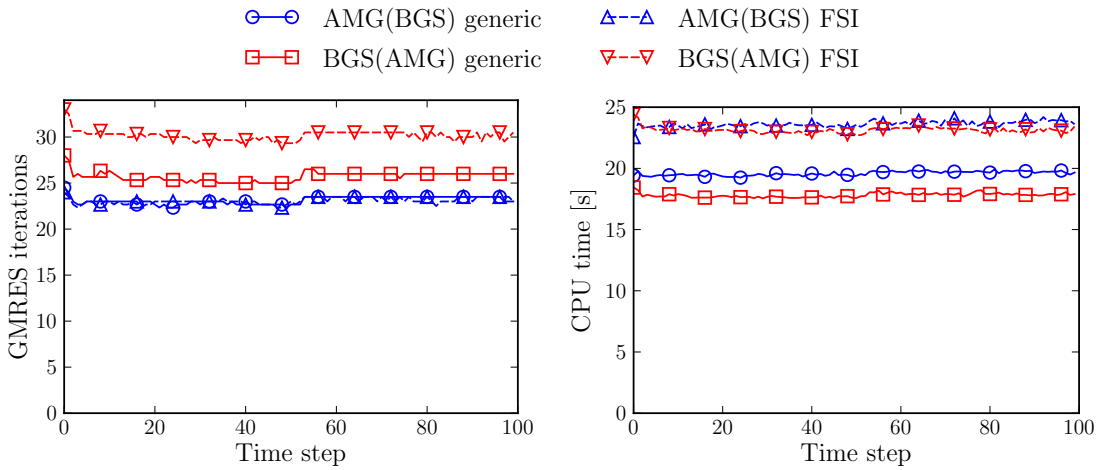

Figure 18: FSI example: Performance of the generic and FSI-specific preconditioners. The CPU times include the setup cost of the preconditioner. The results are for the finest discretization given in Table 10 (2.8 million unknowns on 64 processors).

The same comparison is carried out using a weak scalability study. To this end, four meshes are considered (see Table 10) and the number of processors are chosen such that the amount of work per processor is fairly constant (circa 45000 unknowns per processor). The results (Figure 19) show that the studied methods have a good scalability also for this FSI example. Again, the observed scalabilities do not correspond to the optimal multigrid behavior expected in elliptic problems, but they can be considered decent results for this complex application. Note also that both the generic and the FSI-specific implementations lead to similar CPU times, which confirms that the general purpose preconditioners are efficient methods.

7.4 Patient-specific lung example

Finally, we study the applicability of the proposed methods to the lung model proposed by Yoshihara et al. [23, 24] (see Section 6.3) using a real-world patient-specific example. The problem geometry corresponds to the airway tree and the lung lobes (see Figure 20) of a healthy male subject, which is obtained

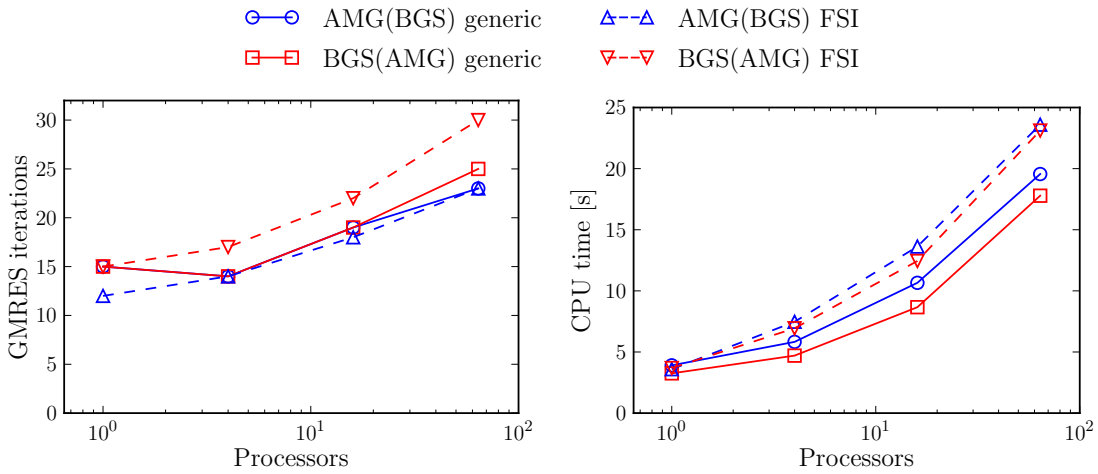

Figure 19: FSI example: Results of the weak scalability study. The CPU times include the setup costs of the preconditioner.

from computer tomography (CT) images provided by the Medical Office "Diagnostische Radiologie" in Stuttgart, Germany. The geometry of the air tree is reconstructed up to the 3rd generation leading to 5 artificial outlets that have to be coupled to the lung tissue. This is achieved by subdividing the lung lobes into several sub-domains associated with each one of the ending branches (see Figure 20). These regions represent the volumes enclosed by the boundaries $\Gamma_\bullet$ in Section 6.3.

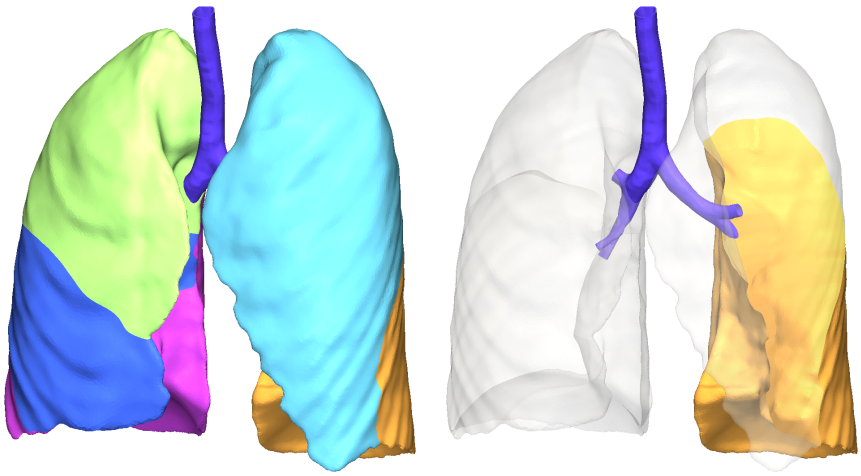

Figure 20: Patient-specific lung example: Lung geometry with the airway tree reconstructed up-to the 3rd generation. The left picture shows the different regions in which the lung is sub-divided according to the model by Yoshihara et al. [23, 24]. The right picture shows the airway tree and one of the these regions.

The lung tissue is modeled with a Neo-Hookean solid with young's modulus $E = 6.75$ kPa, Poisson's ration $\nu = 0.3$ and density $\rho^{\mathrm{s}} = 1.2 \cdot 10^3$ Kg/m^3 which are determined experimentally in [65]. The same material is considered for the airway tree but with a stiffer Young's modulus $E = 13.5$ kPa to mimic the stiffer airway walls. The air is modeled as an incompressible Newtonian fluid with density $\rho^{\mathrm{f}} = 1$ Kg/m^3 and viscosity $\mu = 1.5 \cdot 10^{-5}$ Pa $\cdot$ s. The entire system is assumed to be initially at rest and the action of the thoracic muscles causing the inflation of the lungs is modeled as a displacement boundary condition on the outer surface of the lung lobes. This boundary data is known experimentally form a temporal sequence of CT images. As a result, the lung inflates and causes the suction of the air through the airway tree, see Figure 21.

For the discretization of the problem, we consider linear tetrahedral elements for the structural domain

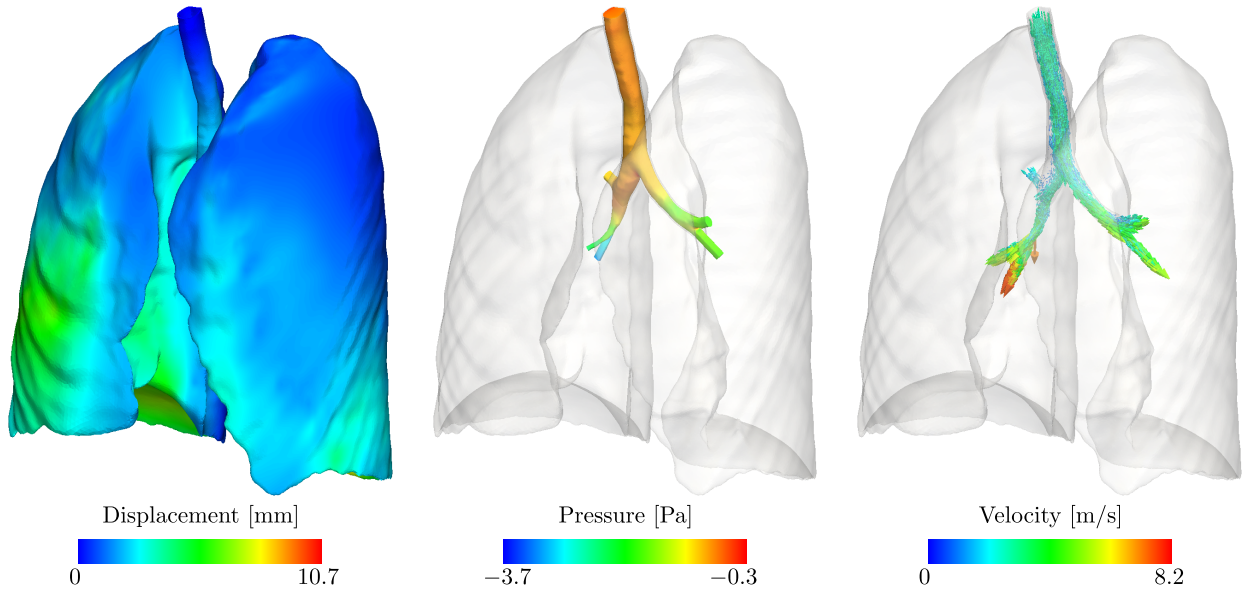

Figure 21: Example 2: Numerical solution of the lung model consisting in the structural displacement (left), fluid pressure (center) and fluid velocity (right) at time $t = 0.75$ s.

(lung lobes and walls of the air tree) and linear tetrahedral stabilized finite elements for the air. The number of degrees of freedom for each field is given in Table 12. The time step is chosen as $\Delta t = 1 \cdot 10^{-3}$s and the problem is simulated for 750 time steps. Convergence of the Newton method is declared when the full non-linear residual fulfills $\|\mathbf{f}^{\mathrm{C;FSI}}\|_{\mathrm{rms}} < 10^{-5}$. The stopping criterion for the GMRES iteration is $\|\mathbf{r}\|_2/\|\mathbf{r}^0\|_2 < 10^{-6}$. For the resolution of this challenging problem we consider the preconditioners SIMPLE(BGS(AMG)) and SIMPLE(AMB(BGS)) presented in Table 3.

Processors	Solid	Fluid	Grid	Constraint	Total	Total/Processors
64	1557897	429552	210498	5	2197952	34343

Table 12: Patient-specific lung example: Number of processors and number of degrees of for the mesh considered in this example.

Figure 22 shows the performance of the studied preconditioners. Note that both methods are able to solve the linear systems in about 20 iterations, which is a very good convergence behavior. This shows that the proposed methods are preconditioning the problem effectively also in this challenging real-world example. The average time spent in the solution of the linear systems is 14.9 seconds for SIMPLE(BGS(AMG)) and 16.6 seconds for SIMPLE(AMG(BGS)). These times represent the 33.4 % of the total simulation time for the SIMPLE(BGS(AMG)) method and 36.5 % for the SIMPLE(AMG(BGS)). These can be considered very good results since the linear solver phase (which is usually considered the bottle neck of the computation) is taking here less than the 40% of the total computation time. In conclusion, the results show that the presented methods are effective preconditioners allowing to solve this challenging problem in decent amount of time.

8 Concluding remarks

Several algebraic multigrid preconditioners have been evaluated for solving challenging systems of linear equations resulting from the discretization of coupled problems. These methods have been considered because they can be implemented for a generic coupled problem and used in different applications. The versatility of the selected techniques is demonstrated by considering three very different coupled

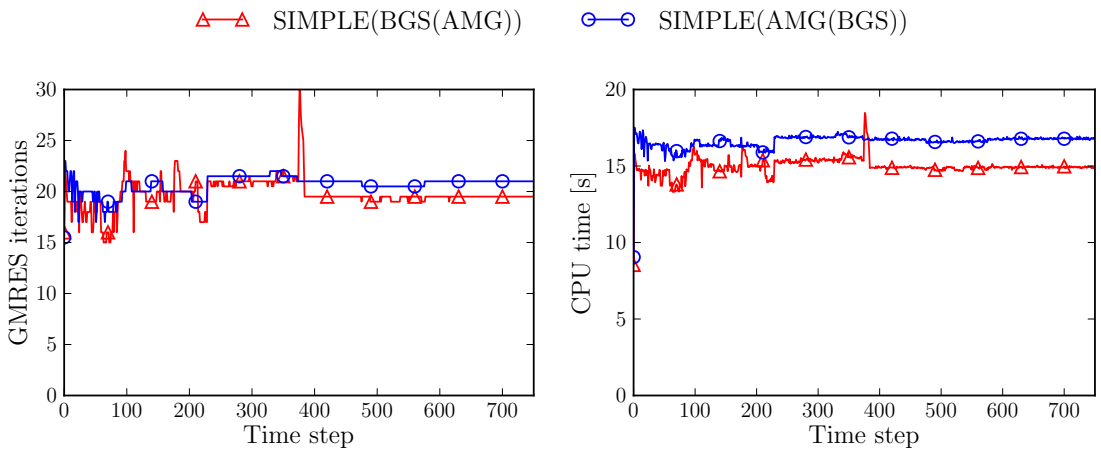

Figure 22: Patient-specific lung example: Average of linear solver iterations per Newton step (left) and the corresponding CPU times (right) including the setup cost of the preconditioners.

problems: thermo-structure interaction (TSI), fluid-structure (FSI) interaction and a complex model of the human lung. The numerical results show that the proposed methods are scalable preconditioners in different settings including a benchmark test for TSI, a real world TSI example, and a FSI problem. Moreover, it is shown that the generic methods are efficient in terms of CPU time even though they are not optimized for a particular problem. This is showed in an FSI example by comparing the general purpose methods presented here and FSI specific preconditioners previously presented in the literature. Finally, the robustness and the wide applicability range of the methods are illustrated with a real-world patient-specific simulation of the human lung. The proposed methods lead to a very good convergence of the preconditioned GMRES method and efficient CPU times also in this challenging setting.

Acknowledgment

We gratefully acknowledge the support through the German Research Foundation (Deutsche Forschungsgemeinschaft – DFG) in the framework of the Sonderforschungsbereich Transregio 40, TP D4.

The authors also acknowledge the support given by the Bayerische Kompetenznetzwerk für Technisch-Wissenschaftliches Hoch- und Höchstleistungsrechnen (KONWIHR) in the framework of the project *Efficient solvers for coupled problems in respiratory mechanics.*

The access to the high performance facilities provided by the Leibniz Supercomuting Center (Munich, Germany) in the framework of the project *Coupled Problems in Computational Modeling of the Respiratory System* (project code pr32ne) is also gratefully acknowledged.

References

[1] C. A. Felippa, K. C. Park, and C. Farhat. Partitioned analysis of coupled mechanical systems. *Comput. Methods Appl. Mech. Engrg.*, 190:3247–3270, 2001. doi:10.1016/S0045-7825(00)00391-1.

[2] P. Causin, J. F. Gerbeau, and F. Nobile. Added-mass effect in the design of partitioned algorithms for fluid-structure problems. *Comput. Methods Appl. Mech. Engrg.*, 194:4506–4527, 2005. doi:10.1016/j.cma.2004.12.005.

[3] C. Förster, W. A. Wall, and E. Ramm. Artificial added mass instabilities in sequential staggered coupling of nonlinear structures and incompressible viscous flows. *Comput. Methods Appl. Mech. Engrg.*, 196:1278–1293, 2007. doi:10.1016/j.cma.2006.09.002.

[4] S. Badia, A. Quaini, and A. Quarteroni. Modular vs. non-modular preconditioners for fluid-structure systems with large added-mass effect. *Comput. Methods Appl. Mech. Engrg.*, 197:4216–4232, 2008. doi:10.1016/j.cma.2008.04.018.

[5] S. Badia, A. F. Martín, and R. Planas. Block recursive LU preconditioners for the thermally coupled incompressible inductionless MHD problem. *J. Comput. Phys.*, 274:562–591, 2014. doi:10.1016/j.jcp.2014.06.028.

[6] E. C. Cyr, J. N. Shadid, R. S. Tuminaro, R. P. Pawlowski, and L. Chacón. A new approximate block factorization preconditioner for two-dimensional incompressible (reduced) resistive MHD. *SIAM J. Sci. Comput.*, 35:B701–B730, 2013. doi:10.1137/12088879X.

[7] C. Danowski, V. Gravemeier, L. Yoshihara, and W. A. Wall. A monolithic computational approach to thermo-structure interaction. *Int. J. Numer. Meth. Engng.*, 95:1053–1078, 2013. doi:10.1002/nme.4530.

[8] M. W. Gee, U. Küttler, and W. A. Wall. Truly monolithic algebraic multigrid for fluid-structure interaction. *Int. J. Numer. Meth. Engng.*, 85:987–1016, 2011. doi:10.1002/nme.3001.

[9] M. Heil, A. L. Hazel, and J. Boyle. Solvers for large-displacement fluid-structure interaction problems: segregated versus monolithic approaches. *Comput. Mech.*, 43:91–101, 2008. doi:10.1007/s00466-008-0270-6.

[10] V. E. Howle and R. C. Kirby. Block preconditioners for finite element discretization of incompressible flow with thermal convection. *Numer. Linear Algebra Appl.*, 19:427–440, 2012. doi:10.1002/nla.1814.

[11] V. E. Howle, R. C. Kirby, and G. Dillon. Block preconditioners for coupled physics problems. *SIAM J. Sci. Comput.*, 35:S368–S385, 2013. doi:10.1137/120883086.

[12] U. Küttler, M. Gee, Ch. Förster, A Comerford, and W. A. Wall. Coupling strategies for biomedical fluid-structure interaction problems. *Int. J. Numer. Meth. Biomed. Engng.*, 26:305–321, 2010. doi:10.1002/cnm.1281.

[13] P. T. Lin, J. N. Shadid, R. S. Tuminaro, M. Sala, G. L. Hennigan, and R. P. Pawlowski. A parallel fully coupled algebraic multilevel preconditioner applied to multiphysics PDE applications: Drift-diffusion, flow/transport/reaction, resistive MHD. *Int. J. Numer. Meth. Fluids*, 64:1148–1179, 2010. doi:10.1002/fld.2402.

[14] Y. Saad and H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7:856–869, 1986. doi:10.1137/0907058.

[15] W. L. Briggs, V. E. Henson, and S. F. McCormick. *A Multigrid Tutorial: Second Edition.* Society for Industrial and Applied Mathematics (SIAM), 2000. ISBN 9780898714623. doi:10.1137/1.9780898719505.

[16] U. Trottenberg, C. W. Oosterlee, and A. Schüller. *Multigrid.* Academic Press, 2001. ISBN 9780127010700.

[17] P. Vaněk, J. Mandel, and M. Brezina. Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems. *Computing*, 56:179–196, 1996.

[18] H. Elman, V. E. Howle, J. Shadid, R. Shuttleworth, and R. Tuminaro. A taxonomy and comparison of parallel block multi-level preconditioners for the incompressible Navier-Stokes equations. *J. Comput. Phys.*, 227:1790–1808, 2008. doi:10.1016/j.jcp.2007.09.026.

[19] L. S. Caretto, A. D. Gosman, S. V. Patankar, and D. B. Spalding. Two calculation procedures for steady, three-dimensional flows with recirculation. In *Proceedings of the Third International Conference on Numerical Methods in Fluid Mechanics*, volume 19, pages 60–68. 1973. doi:10.1007/BFb0112677.

[20] E. C. Cyr, J. N. Shadid, and R. S. Tuminaro. Stabilization and scalable block preconditioning for the Navier–Stokes equations. *J. Comput. Phys.*, 231:345–363, 2012. doi:10.1016/j.jcp.2011.09.001.

[21] E. G. Phillips, H. C. Elman, E. C. Cyr, J. N. Shadid, and R. P. Pawlowski. A block preconditioner for an exact penalty formulation for stationary MHD. *SIAM J. Sci. Comput.*, 36:B930–B951, 2014. doi:10.1137/140955082.

[22] W. A. Wall and M. W. Gee. BACI a multiphysics simulation environment. Technical report, Technische Universität München, 2014.

[23] L. Yoshihara, M. Ismail, and W. A. Wall. Bridging scales in respiratory mechanics. In Gerhard A. Holzapfel and Ellen Kuhl, editors, *Computer Models in Biomechanics*, pages 395–407. Springer Netherlands, 2013. doi:10.1007/978-94-007-5464-5_28.

[24] L. Yoshihara, C. J. Roth, and W. A. Wall. Fluid-structure interaction including volume coupling of homogenized subdomains for modeling respiratory mechanics. *Submited*, 2015.

[25] E. Chow and M. A. Heroux. An object-oriented framework for block preconditioning. *ACM Trans. Math. Softw.*, 24(2):159–183, 1998. doi:10.1145/290200.287639.

[26] V. E. Howle, R. C. Kirby, K. Long, B. Brennan, and K. Kennedy. Playa: High-performance programmable linear algebra. *Scientific Programming*, 20(3):257–273, 2012. doi:10.3233/SPR-2012-0347.

[27] E. C. Cyr, J. N. Shadid, and R. S. Tuminaro. Teko: A block preconditioning capability with concrete example applications in Navier-Stokes and MHD. *SIAM Journal on Scientific Computing*, 2016. Accepted for publication.

[28] TEKO package home page. `https://trilinos.org/packages/teko/`, 2016. [Online accessed 30-May-2016].

[29] J. Brown, M. G. Knepley, D. A. May, L. C. McInnes, and B. Smith. Composable linear solvers for multiphysics. In *Proceedings of the 2012 11th International Symposium on Parallel and Distributed Computing*, pages 55–62. IEEE Computer Society, 2012. doi:10.1109/ISPDC.2012.16.

[30] PCFIELDSPLIT documentation web page. `http://www.mcs.anl.gov/petsc/petsc-current/docs/manualpages/PC/PCFIELDSPLIT.html`, 2016. [Online accessed 30-May-2016].

[31] The Trilinos Project web page. `https://trilinos.org/`, 2016. [Online accessed 30-May-2016].

[32] PETSc home page. `https://www.mcs.anl.gov/petsc/`, 2016. [Online accessed 30-May-2016].

[33] H. A. van der Vorst. *Iterative Krylov Methods for Large Linear Systems*. Cambridge Monographs on Applied and Computational Mechanics. Cambridge University Press, 2003. ISBN 0-521-81828-1.

[34] T. Ginsburg. The conjugate gradient method. *Numer. Math.*, 5:191–200, 1963. doi:10.1007/BF01385890.

[35] M. A. Heroux. AztecOO user guide. Technical Report SAND2004-3796, Sandia National Laboratories, 2007. URL `http://trilinos.org/packages/aztecoo/`.

[36] Bau D. Trefethen, L. N. *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, 1997.

[37] Y. Saad. *Iterative Methods for Sparse Linear Systems: Second Edition*. Society for Industrial and Applied Mathematics, 2003. ISBN 9780898715347.

[38] J. Liesen and P. Tichy. Convergence analysis of krylov subspace methods. *GAMM-Mitteilungen*, 27: 153–173, 2004. doi:10.1002/gamm.201490008.

[39] P. Wesseling and C. W. Oosterlee. Geometric multigrid with applications to computational fluid dynamics. *J. Comput. Appl. Math.*, 128:311–334, 2001. doi:10.1016/S0377-0427(00)00517-3.

[40] P. Vaněk, M. Brezina, and J. Mandel. Convergence of algebraic multigrid based on smoothed aggregation. *Numer. Math.*, 88:559–579, 2001. doi:10.1007/s211-001-8015-y.

[41] A. Prokopenko, J. J. Hu, T. A. Wiesner, C. M. Siefert, and R. S. Tuminaro. Muelu user's guide 1.0. Technical Report SAND2014-18874, Sandia National Labs, 2014. URL `http://trilinos.org/packages/muelu/`.

[42] M. W. Gee, C. M. Siefert, J. J. Hu, R. S. Tuminaro, and M. G. Sala. ML 5.0 smoothed aggregation user's guide. Technical Report SAND2006-2649, Sandia National Laboratories, 2006. URL `http://trilinos.org/packages/ml/`.

[43] S. C. Brenner and R. Scott. *The Mathematical Theory of Finite Element Methods*. Texts in Applied Mathematics. Springer, 2008. ISBN 9780387759333.

[44] P. T. Lin, M. Sala, J. N. Shadid, and R. S. Tuminaro. Performance of fully coupled algebraic multilevel domain decomposition preconditioners for incompressible flow and transport. *Int. J. Numer. Meth. Engng.*, 67:208–225, 2006. doi:10.1002/nme.1624.

[45] P. T. Lin, J. N. Shadid, M. Sala, R. S. Tuminaro, G. L. Hennigan, and R. J. Hoekstra. Performance of a parallel algebraic multilevel preconditioner for stabilized finite element semiconductor device modeling. *J. Comput. Phys.*, 228:6250–6267, 2009. doi:10.1016/j.jcp.2009.05.024.

[46] O. Axelsson and M. Neytcheva. A general approach to analyse preconditioners for two-by-two block matrices. *Numer. Linear Algebra Appl.*, 20:723–742, 2013. doi:10.1002/nla.830.

[47] M. Sala and M. Heroux. Robust algebraic preconditioners with IFPACK 3.0. Technical Report SAND-0662, Sandia National Laboratories, 2005. URL `https://trilinos.org/packages/ifpack/`.

[48] M. Sala, K. Stanley, and M. Heroux. Amesos: A set of general interfaces to sparse direct solver libraries. In *Proceedings of PARA'06 Conference, Umea, Sweden*, 2006.

[49] J. P. Van Doormaal and G. D. Raithby. Enhancements of the simple method for predicting incompressible fluid flows. *Numer. Heat Transfer*, 7:147–163, 1984. doi:10.1080/01495728408961817.

[50] M. Mayr, T. Klöppel, A. W. Wall, and M. W. Gee. A temporal consistent monolithic approach to fluid-structure interaction enabling single field predictors. *SIAM J. Sci. Comput.*, 37:B30–B59, 2015. doi:10.1137/140953253.

[51] F. Armero and J. C. Simo. A new unconditionally stable fractional step method for nonlinear coupled thermomechanical problems. *Int. J. Numer. Meth. Engng.*, 35:737–766, 1992. doi:10.1002/nme.1620350408.

[52] P. Erbts and A. Düster. Accelerated staggered coupling schemes for problems of thermoelasticity at finite strains. *Comput. Math. Appl.*, 64:2408–2430, 2012. doi:10.1016/j.camwa.2012.05.010.

[53] C. Farhat, K. C. Park, and Y. Dubois-Pelerin. An unconditionally stable staggered algorithm for transient finite element analysis of coupled thermoelastic problems. *Comput. Methods Appl. Mech. Engrg.*, 85:349–365, 1991. doi:10.1016/0045-7825(91)90102-C.

[54] C. Miehe. Entropic thermoelasticity at finite strains. Aspects of the formulation and numerical implementation. *Comput. Methods Appl. Mech. Engrg.*, 120:243–269, 1995. doi:10.1016/0045-7825(94)00057-T.

[55] U. Küttler and W. A. Wall. Fixed-point fluid-structure interaction solvers with dynamic relaxation. *Comput. Mech.*, 43:61–72, 2008. doi:10.1007/s00466-008-0255-5.

[56] C. Förster, W. A. Wall, and E. Ramm. On the geometric conservation law in transient flow calculations on deforming domains. *Int. J. Numer. Meth. Fluids*, 50:1369–1379, 2006. doi:10.1002/fld.1093.

[57] P. Crosetto, S. Deparis, G. Fourestey, and A. Quarteroni. Parallel algorithms for fluid-structure interaction problems in haemodynamics. *SIAM J. Sci. Comput.*, 33:1598–1622, 2011. doi:10.1137/090772836.

[58] M. Heil. An efficient solver for the fully coupled solution of large-displacement fluid-structure interaction problems. *Comput. Methods Appl. Mech. Engrg.*, 193:1–23, 2004. doi:10.1016/j.cma.2003.09.006.

[59] V. I. Danilovskaya. On a dynamical problem of thermoelasticity. *Prikladnaya Matematika i Mekhanika*, 16:341–344, 1952.

[60] M. Tanaka, T. Matsumoto, and M. Moradi. Application of boundary element method to 3-d problems of coupled thermoelasticity. *Eng. Anal. Bound. Elem.*, 16:297–303, 1995. doi:10.1016/0955-7997(95)00074-7.

[61] M. Sala, K. S. Stanley, and M. A. Heroux. On the design of interfaces to sparse direct solvers. *ACM Trans. Math. Softw.*, 34:9:1–9:22, 2008. doi:10.1145/1326548.1326551.

[62] V. K. Arya and S. M. Arnold. Viscoplastic analysis of an experimental cylindrical thrust chamber liner. *AIAA Journal*, 30:781–789, 1992. doi:10.2514/3.10985.

[63] J.-F. Gerbeau and M. Vidrascu. A quasi-newton algorithm based on a reduced model for fluid-structure interaction problems in blood flows. *ESAIM: Math. Model. Num.*, 37:631–647, 2003. doi:10.1051/m2an:2003049.

[64] M. Sala and R. Tuminaro. A new Petrov-Galerkin smoothed aggregation preconditioner for non-symmetric linear systems. *SIAM J. Sci. Comput.*, 31:143–166, 2008. doi:10.1137/060659545.

[65] S. M. K. Rausch, C. Martin, P. B. Bornemann, S. Uhlig, and W. A. Wall. Material model of lung parenchyma based on living precision-cut lung slice testing. *J. Mech. Behav. Biomed. Mater*, 4: 583–592, 2011. doi:10.1016/j.jmbbm.2011.01.006.